Oracle® Communications IP Service Activator

User's Guide

Release 7.2

E35290-01

October 2013

ORACLE®

Oracle Communications IP Service Activator User's Guide, Release 7.2

E35290-01

Contents

3 Defining and Applying Roles

4 Setting Up Domain Information

5 Discovering and Setting Up the Network

6 Representing and Mapping Objects

7 Checking Device Status and Capabilities

8 Working with Transactions

Preface

This guide provides detailed step-by-step information about setting up the Oracle Communications IP Service Activator system as well as discovering and representing the network topology.

Audience

This document is intended for network configuration engineers responsible for the initial setup of IP Service Activator after installation.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

1

IP Service Activator Setup: Overview

This chapter outlines the steps to install and run Oracle Communications IP Service Activator. For more detailed information about installing IP Service Activator, see *IP Service Activator Installation Guide*.

Customizing IP Service Activator

Customizing IP Service Activator for your network and applying services and policies to it requires the following steps:

1. Install the system.
2. Configure the system.
3. Set up global (system and domain) information.
4. Define policy roles and role assignment rules.
5. Discover the network and set up device information.
6. Set up VPN services.
7. Set up basic policy data.
8. Set up QoS or security policies.
9. Set up SLA monitoring.

Installing IP Service Activator

In a normal distributed operation, IP Service Activator components are installed on multiple hosts. For tests and evaluations, we recommend that you install a single host system. This allows you to become familiar with IP Service Activator before installing a fully-operational distributed system. We also recommend that the test environment represent the live deployment, based on the performance guidelines.

For full details of the installation process, see *IP Service Activator Installation Guide*.

Configuring IP Service Activator

After installation of IP Service Activator, there are a number of configuration tasks that must be performed. Some of the aspects that must be configured include setting the IP addresses and ports for the different components of IP Service Activator, specifying connectivity information for the database, specifying which components of IP Service Activator to run, and setting tuning parameters for the various components.

For full details of the installation process, see *IP Service Activator System Administrator's Guide*.

Setting Up System and Domain Information

There are a number of tasks that you should do immediately after installation. These include the following:

- Setting up system users and groups: one initial system user is created the first time you log in. You need to set up security options for that user and create additional user groups and users.
- Setting up domains: you can set up multiple domains. Autonomous System (AS) regions can be configured globally, and multiple AS regions can be configured on a domain. Device AS regions override global AS configurations.
- Loading basic configuration data: you can install set-up files for each domain that automatically create the basic data that you need when setting up services and policies.

These tasks are described in detail in *IP Service Activator System Administrator's Guide*.

Setting Up Roles

In IP Service Activator, roles provide the meeting point between policy elements and policy targets, such as devices, interfaces, and sub-interfaces. The roles you need to set up depend on the service or policy you plan to implement:

- IP Service Activator provides a set of system-defined roles. In order to configure MPLS VPNs, you must apply these roles to the devices and interfaces that are to be managed.
- User-defined roles can be created and assigned to policy targets.

See "Defining and Applying Roles" for more information about defining roles.

Discovering the Network

The next step is to set up details of the network for each domain that you are managing. You do this by running a device discovery process that uses SNMP to find out information about devices and connected segments.

The discovery process does the following:

- Finds out details of the devices, segments, and hosts in the network
- Assigns devices to the proxy agents that will manage them
- Sets up the security parameters that IP Service Activator needs to configure devices
- Ascertains the capabilities of devices and interfaces, indicating the QoS, access control, measurement, and VPN features that are available

Before running the discovery process you must set up certain device-specific information, and after discovery is complete you must set up IP Service Activator to manage the devices.

See "Discovering and Setting Up the Network" for full details about discovering devices and the tasks you must perform before and after device discovery.

Mapping the Network

The topology map provides a visual representation of the discovered network.

If IP Service Activator is set to map the network automatically, nodes are added to the topology map as they are discovered. By default, discovered nodes must be mapped manually when the discovery process is complete.

See "Representing and Mapping Objects" for information about mapping the network.

Setting Up VPN Services

After discovering the devices and network interfaces, creating new network interfaces as required, and mapping the network, you can create the VPN services that you want to implement in each domain:

- Multi-Protocol Label Switching VPNs (MPLS VPNs): you set up a VPN by defining customers and sites and specifying how sites are linked together.
- Transparent LAN Service (TLS): you set up a TLS by defining customers and layer 2 sites that specify the criteria on which access to the TLS is based. Layer 2 sites are linked together in a TLS object.
- Circuit Cross Connects (CCCs): you set up a CCC by defining customers and linking VC endpoints. CCCs are specific to Juniper M-series devices.
- Layer 2 Martini VPNs: you set up a Layer 2 Martini VPN by defining customers, creating the Layer 2 Martini VPN and linking Martini endpoints. Layer 2 Martini VPNs are specific to Cisco and Juniper M-series devices.
- VRF-Lite: you set up VRF-Lite by logically partitioning a CE device into groups of interfaces, creating a Virtual CE object, and associating the interfaces from the physical CE to the Virtual CEs.

For more information about VPN services, see *IP Service Activator VPN User's Guide*.

Setting Up Basic Policy Data

If you plan to implement a QoS or security policy, you need to set up the basic data used in configuring these elements.

- Class of service information: defining how differentiated services are being used. You need to set this up if you are setting up a QoS policy.
- Rule components: traffic types, classifications, IP protocols, packet markings, and date and time templates, used when setting up QoS and policy rules.

For information about setting up basic data, see *IP Service Activator QoS User's Guide*.

Setting Up Policies

If you have set up basic policy data, you can create the policies that you want to implement in each domain.

- QoS policies: classification rules, policing rules and PHB groups can be implemented to manage and prioritize categories of traffic at any point in the network.
- Access control policies: access rules can be set up to deny or explicitly permit defined categories of traffic.

For more information about setting up QoS policies, see *IP Service Activator QoS User's Guide*.

Setting Up SLA Monitoring

IP Service Activator integrates with third-party reporting tools to provide SLA monitoring. Various measurement types are available to generate the source data for reporting:

- Service Assurance Agent (SAA) enables you to monitor point-to-point connections in an MPLS VPN or a measurement-only VPN.
- NetFlow and MIB-based measurements enable you to monitor activity at a specific point in the network.

For more information about setting up SLA monitoring, see *IP Service Activator Network and SLA Monitoring Guide*, which provides details of setting up SLA measurement within IP Service Activator.

Transactions in IP Service Activator

IP Service Activator uses a transaction-based model for updating the system and implementing configuration changes. This means that you can make changes through the IP Service Activator client and implement them immediately or save them in a pending state and implement them at another time.

Each transaction has a status on it. Transaction monitor connects to the IP Service Activator Policy Server through the OSS Integration Manager (OIM) and allows upstream systems to confirm that activation transactions have been successfully applied to network elements. See *IP Service Activator System Administrator's Guide* for more information.

The tasks associated with initial system setup are generally performed by a single user as a one-time task. You can implement these changes immediately. For implementing VPN services and setting up policies, you can break down configuration changes into a number of transactions and perform a phased implementation.

See "Working with Transactions" for more information about transactions.

2

The IP Service Activator Client

This chapter introduces elements of Oracle Communications IP Service Activator client and describes how to navigate the system and customize the interface. It covers the following areas:

- Appearance of the client when you run the system for the first time
- Options available for issuing commands: menus, toolbars, and context menus
- Window types used when configuring the system
- How the system models the network applied policies in IP Service Activator object model
- Available navigation tools
- How to change the appearance of the client
- How to search IP Service Activator, including searching for concrete configuration
- How to print

Introduction

The client provides a representation of the service that is managed by IP Service Activator. The client acts as a front end to the policy server, the core component of IP Service Activator.

Working within the client, you can discover the devices to be managed, map the topology of the network and construct the policy to be applied. The information you enter in the client is saved to a database via the policy server.

IP Service Activator models the network and the policy you apply to it using an object model, where every element is an object with its own properties and capabilities. Every object type has an icon associated with it. Where the state of an object may change, color is used to represent the state of the object. See "Status Context View" for more information.

The system provides a policy view for objects that can have a policy applied to them, where object icons are displayed in shades of gray that reflect their policy role.

You can have several windows onto the system open simultaneously, with each one offering the same or a different view of the system. These windows are known as domain management windows.

Each window is divided into several panes. Each pane displays a different type of information. You can choose whether to hide or show selected panes.

There may be several client components running at any one time, with each user entering new configuration details. As a user makes changes in the client, the changes are queued locally, and are not reflected in remote clients. When a user commits a transaction, the policy server co-ordinates the information between clients, ensuring that all users' views of the system are consistent.

Running IP Service Activator for the First Time

When you run IP Service Activator for the first time after installation, there are no domains defined and only default system information is displayed on the **System** tab.

The first time the client opens, the system prompts you to change the administrator password, which has a default value of **admin**. Changing the password is a mandatory step.

The main IP Service Activator screen is referred to as the Global Setup window and is always displayed when you run IP Service Activator. See "The Global Setup Window" for more information.

User Access Levels and the Client

The appearance of the client is affected by your security level. You may be able to see all or only parts of the system. For example, you may be unable to see the **System** tab or some menu items. Similarly you may be able to see some objects, but unable to modify their values. For example, you may be able to see a list of user-defined roles but unable to modify them or add new roles to the list.

This guide, and the online Help system, assume you can view and modify all parts of the system. This access level is available to users with Super User access. If you have a lower access level you may be unable to access some parts of IP Service Activator described in this guide.

Note that the **Policy** tab is not displayed to the user and options for creating rules and PHB groups are missing from the device's context menu. The user is therefore unable to view or create policies.

Note also that you cannot see faults associated with object types to which you do not have access. In this example, the warnings that relate to a PHB group with no roles associated with it cannot be seen in the second illustration.

IP Service Activator Windows

Within the IP Service Activator client you can open multiple windows and arrange them as needed to support whatever you are currently working on. For example, you can display windows side by side to compare information between two domains, or drag and drop objects between windows on the same or different domains.

Information is displayed in the following window types:

- Domain management windows display information about a particular policy domain. This window type is your workspace area for managing the domain.
- The Global Setup window displays global information and allows you to create and manage multiple domains.

The information that is displayed in each window type and the actions you can perform depend on your access level. For example, in the Global Setup window you may be able to view existing domains but unable to create new domains. See "User Access Levels and the Client" for more information.

To open a new domain management window:

1. In a domain management window, either:

   - From the **Window** menu, select **New Window**.

     A new domain management window opens. The window's Hierarchy pane shows the complete object model for the domain.

   or:

   - With an object selected in the Hierarchy pane, from the **View** menu, select **Open**.

     A new domain management window opens. The window's Hierarchy pane shows only those objects that exist below the object you selected in the Hierarchy pane.

You can also create a new window by double-clicking on a domain on the **Domains** tab in a Global Setup window.

The Domain Management Windows

Domain management windows provide a workspace for configuring a policy domain. You work in a domain management window to:

- Discover the network you intend to manage
- Create and view topology maps of the network
- Set up the policies and services you want to implement
- Create transactions and send configuration details to the devices you are managing

If you are using the Network Processor, you must have all the devices present in the client in order to create capabilities, options, and the MIPSA registry. For information about using the Network Processor, see *IP Service Activator Cisco IOS Cartridge Guide*.

The following shows a domain management window with the topology map displayed:

When you run IP Service Activator for the first time, there are no domains created and you cannot display a domain management window. See "Setting Up Domains" for information about creating a domain.

Panes in the Domain Management Window

The domain management window is divided into sections.

See "Hiding and Showing Windows and Window Elements" for information about hiding and showing panes in the domain management window.

Hierarchy Pane and Tabs

The Hierarchy pane displays the IP Service Activator object model as a hierarchical tree structure, showing the major container (parent/child) relationships between the objects. Associated objects are arranged into tabbed groups. You can access all of the objects within the domain by selecting tabs and browsing through the hierarchy.

Branches of the tree can be expanded or collapsed by clicking the '+' or '-' symbols in the tree.

By default, the hierarchy tree appears at the top left of each explorer window, but it can be moved, deleted, or resized using the normal window controls.

Tabs allow you to view different groupings of data. By default, the tabs appear at the left of the Hierarchy pane. You can change their position by clicking in the title bar of the hierarchy tree.

See "Tabs in the Domain Management Window" for more information about the hierarchy tree tabs. See "How IP Service Activator Models the System" for information about the IP Service Activator object model.

Details Pane

The Details pane displays information about the object that you have selected in the Hierarchy pane, and the network topology around that object.

Three types of views are available:

- **Report view** displays the contents in the form of a list of objects, together with additional columns of attributes
- **Compact List view** displays the contents as a list of objects
- **Map view** displays the network topology in graphical form. You can only select the map view when a network or VPN object is selected

Information is displayed in a map, list or report view, depending on the selected object's type. You can double-click on an object in the details pane to drill down to a lower level. If there are no child objects, the Properties dialog box for the selected object is displayed.

For example:

- If you have selected a network object, its details can be viewed as a topology map, a list of the devices within that network or a report list showing details of each device's state, role, IP address, and so on.
- If you have selected a VPN object, the sites within the VPN can be viewed in map, list, or report form.
- If you have selected a device, the interfaces available on the device can be displayed in list or report form.

You can also display information about the configuration that applies to an object in the Details pane or the content of any of the system's log files. For more information, see *IP Service Activator QoS User's Guide* and *IP Service Activator VPN User's Guide*. For information about IP Service Activator log files, see *IP Service Activator System Administrator's Guide*.

See "Printing" for information about printing the content of the Details pane.

Ancestry Pane

The Ancestry pane displays the object model as a reverse tree, with the object that you have currently selected in the Hierarchy pane shown at the top level and its parent objects beneath it. This reverse hierarchy is useful for tracing all the links between an object and its parent objects. Note that where links are created, an object can have several parents.

The ancestry tree is not a navigation tool, that is, selecting items in the tree does not affect the information that is displayed in the Details pane. See "How IP Service Activator Models the System" for more information about the object model.

Current Faults Pane

The Current Faults pane displays information, warning, and error messages reported by the system, as shown in Figure 2–1. Most messages result from network discovery or data validation.

Figure 2–1 The Current Faults Pane

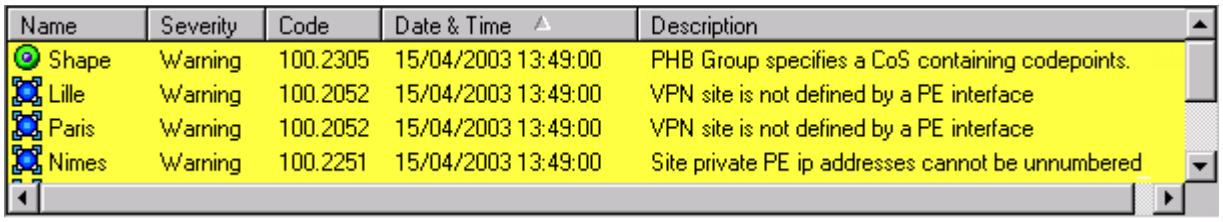

The severity, code number, date and time of occurrence, and error description appear. Each message type is color-coded for reference.

In most cases, after a fault is fixed, its corresponding message is automatically removed. If the fault is not automatically removed, you must delete the fault manually, which triggers IP Service Activator to push the data from the Policy Server to the corresponding driver.

To display more information about a particular message, select it and press F1 or choose **Help** from the context menu.

By default the pane appears at the bottom of the main window, but you can detach it and move it to a different position on your desktop. Click the frame surrounding the window and drag it to a new position.

When you restart the IP Service Activator client, it opens in the domain view.

Current Faults Message Content

Messages in the Current Faults pane are categorized according to their degree of criticality. Table 2–1 describes fault categories and what those faults mean.

Table 2–1 Current Faults Message Content

Category	Fault Type
Messages on a red background	Indicate critical faults requiring intervention, such as failure of an IP Service Activator system component or a serious problem with a device.
Messages on an orange background	Indicate errors, i.e. serious faults in the policy system, but not in a system component.
Messages on a yellow background	Indicate warnings, such as validation errors.
Messages on a blue background	Indicate notices. Some user action is required.
Messages on a cyan background	Indicate information only; no user action required.

For each message, the following information appears:

- Name: Name of the object in which the fault occurred. Note that the icon of the object also appears.
- Severity: Critical, Error, Warning, Notice, or Information, as described in Table 2–1.

- Code: Code number of the message. Numbers prefixed with 100 indicate those generated from Service Activator components. Any other prefixes indicates a message originating from a component not supplied by Service Activator, for example, device drivers produced by a third party. For more information about these messages, please contact your software supplier.
- Date & Time: Date and time that the error occurred.
- Description: Text of the message.

For a list of messages that can occur, and suggestions about how to fix the underlying problem, see IP Service Activator online Help.

To display information about a particular message, select the message in the Current Faults pane and press F1 or choose **Help** from the context menu.

By default the Current Faults pane appears at the bottom of the main window. However, it is dockable at the top, bottom, left, or right of the main application window. Click the frame surrounding the window and drag it to the new position. Alternatively, if you drag it outside the main application window, it becomes a separate, floating window.

A fault message appears in the Current Faults pane only while the fault exists; after it is fixed the message is deleted. However all reported faults except validation messages appear in the System Messages report.

You can choose whether or not to display the Current Faults pane. Choose **Current Faults Pane** from the **View** menu.

You can double-click a message, or choose **Properties** from the context menu, to display the property pages for the object on which the fault is reported.

Each fault represents an object and you can drag and drop it in the same way as the object itself. For example, you can link objects by dragging them to an appropriate place on the object model, and if you drag a fault on to the map, the object to which it is related will appear on the map.

Choose **Delete** from the context menu to manually clear the fault. Note that if the problem is not fixed, the fault will re-occur.

Statistics Summary Pane

The Statistics Summary pane summarizes status information for PHB groups, policy rules, CCCs, VPNs, Layer 2 Martini VPNs, and driver scripts. The System Statistics pane can be docked to the top or bottom of the window.

Object statuses include:

- Inactive: the object is created but has not been propagated to the proxy agents
- Disabled: the object is disabled
- Active: the object is propagated to the proxy agents
- Conflict: the object conflicts with an existing object or has failed validation
- Rejected: one or more of the objects failed to install
- Mismatched: the object is missing an installed command from its current configuration
- Installed: an object that is installed on the designated device
- Script Failed: the proxy agent experienced a failure trying to install the driver script and it has therefore been discarded (driver scripts only)

- Totals: the number of objects in each state

Status Bar

The Status Bar displays information about the state of IP Service Activator when you discover devices, propagate configuration data, or retrieve device capabilities, as shown in Figure 2–2.

Figure 2–2 Status Bar

Palette

The Palette, as shown in Figure 2–3, appears when a network map is displayed in Service Activator. In context-sensitive mode, it displays objects relevant to the current selection. For example, if a device is selected, the Palette lists the device's interfaces. If context-sensitivity is turned off, the Palette displays all objects that are not currently displayed on the map. You can display objects on the network map by dragging them from the Palette into the map area, and remove objects from the display by dragging them from the map onto the Palette.

Figure 2–3 The Palette

See "Recalculating a Map Layout" for more information about maps.

Tabs in the Domain Management Window

Within the Domain management window, sets of related objects are grouped and displayed on separate tabs. The tabs are part of the Hierarchy pane. If the Hierarchy pane is not displayed, the tabs are not visible. On each tab, objects are grouped into folders and, in the case of the Topology tab, networks.

You can change the position of the Hierarchy pane's tabs. See "Changing the Hierarchy Pane's Tab Position and Style" for more information.

Objects are grouped under the following tabs:

- **Topology**: network elements, including networks, devices, and interfaces. Work on this tab to discover the network, view and create maps of the network to be managed and assign services and policies to the network elements.
- **Policy**: Class of Service (CoS) and policy building blocks. Work on this tab to create roles, define the levels of service that can be applied to traffic and set up standard and MQC PHB groups, and to define elements such as IP protocols and traffic types. For information about the elements that you can define, see *IP Service Activator QoS User's Guide*.

  Note: You can create some basic rule component data by loading policy configuration files. See "Loading Policy Configuration Data" for more information.

- **Service**: information relating to VPN services. Work on this tab to create and maintain customers, VPNs and sites, TLSs and layer 2 sites, VRF-lite sites, and point-to-point connections.
- **Accounts**: group, subnet, host and user accounts. Work on this tab to create the accounts that act as source or destination points to which rules can be applied.
- **System**: details of the current global parameters and system components. Work on this tab to view information about IP Service Activator system components and their host machines, add and edit user details, access the system's log files, view information about pending, scheduled and committed transactions, define external systems and create SAA templates.

  The System tab includes the following folders:

  - Event Subscriptions
  - External Systems
  - SAA Templates
  - System Hosts - shows the physical relationship between the system components - each of the hosts and the components that have been installed on them. All the host systems on which components are running are listed. These are located and defined by the system and should not be modified.

    Each system host object contains the objects which represent the components that are running on that particular system such as policy server, proxy agent, or system logger.
  - System Logs - lists the logs maintained by the system such as Audit Trail, Device Configuration Log, Event Log and System Messages Log.
  - System User Groups - includes all the users that have been set up in the system. When the system is first installed, one default Super User is created. Additional users should be set up and granted suitable access rights by the system administrator.
  - Transactions
  - SNMP Profiles

  - Policy Server - this object shows the logical relationship between the system components and devices, i.e. the hierarchy of proxy agents, device drivers and the actual devices they control. The policy server object lists the proxy agents that it controls.

- **Custom**: driver scripts that currently exist in IP Service Activator. Work on this tab to import, run and create driver scripts.

The Global Setup Window

The Global Setup window lists the policy domains that are managed by the system and displays general information about each one.

The window has three tabs:

- **Domains**: lists the domains that have been created and displays summary information for each one.
- **System**: this tab is identical to that displayed in the domain management window.
- **Custom**: The custom tab lists the driver scripts that currently exist in the system. Work on this tab to create or import new driver scripts or to view details of existing scripts. The tab includes the Driver scripts folder which holds all driver scripts existing in the system. The custom tab is identical to that displayed in the domain management window.

How IP Service Activator Models the System

IP Service Activator regards anything that it can manipulate as an object. An object may be a representation of a physical element, such as a device or an interface, or a logical element, such as an account or a rule. Information about the objects that exist in the system is held in the database.

Every object has a set of attributes associated with it that define, at the minimum, the name of the object or more complex characteristics, such as detailed information about queuing mechanisms. The number and type of attributes vary between different types of object. For example, a device object has attributes that include its name and IP address, its role, and the interfaces available on the device. A packet marking object, by contrast, has only two attributes associated with it that define its name and the marking associated with it. The combined attributes of an object are referred to as object properties. You can edit an object's properties in its dialog box. To open the properties dialog box for an object, right click on the object and select **Properties**.

Information about objects, the relationships that can exist between them and the actions you can perform on them, together make up the object model. The IP Service Activator object model represents three main areas:

- The underlying network topology
- The policies and services applied to the network
- The IP Service Activator system configuration itself

You can view the objects that are currently defined in the system in the domain management window. You can also view information about system and domain-related objects in the Global Setup window. See "Tabs in the Domain Management Window" for more information.

You can view information about the relationships that exist between the objects in IP Service Activator using the Hierarchy and Ancestry panes:

- The Hierarchy pane shows an object's relationship to objects that exist at a lower level in the object model. You can use the hierarchy tree to drill down through the hierarchy.
- The Ancestry pane shows how an object is related to objects that exist at a higher level in the object model.

You can view the properties that are associated with an object by selecting **Properties** from the object's context menu or from the **View** menu. See "Viewing and Editing Object Properties" for more information.

Modeling the Network, Services, and Policies

As you work in the client, you model the network and the policies or services to be applied to it by:

- Creating objects such as VPNs, classes of service, and sites.
- Defining each object's properties or attributes, such as the connectivity type of a VPN or the name of a site.
- Creating relationships between objects. Objects exist within a hierarchy of parent/child relationships, with the policy, system, or domain objects existing at the highest level. Every new object created is a child of an existing object. You can also make parent/child relationships between some objects by linking them. See "Linking Objects" for more information about linking objects.

IP Service Activator maintains a central version of the object model called the common object model, which is updated when you commit or save a transaction from a remote client. For more information, see *IP Service Activator Concepts*.

Inheritance Between Objects

There is a hierarchy of inheritance between objects, where objects at a lower level in the hierarchy inherit parameters that are applied to objects at a higher level.

Note that devices and interfaces must be tagged with the appropriate roles for policy or parameters to be inherited. The IP Service Activator inheritance model is described in *IP Service Activator QoS User's Guide*.

Working with Objects

You do many of the tasks that are associated with objects using object context menus. The options on each object's menu depend on the object type.

Note: Folders also have context menus associated with them and some object types are created from these menus. See "Creating New Objects" for more information.

To view an object's context menu:

1. Select the object and right-click.

You can not edit information directly within the explorer windows. Each object has a Properties dialog box, which allows you to view details and modify them where necessary. Most Properties dialog boxes contain multiple property pages, selectable from the tree display on the left side of the dialog box. Each property page contains controls for various types of attributes for the object.

The Properties dialog box also appears when you double-click any object in the Details or Hierarchy pane that has no child objects.

Properties dialog boxes are modeless windows; you can leave them open while you select other objects from the client. As different objects are selected, the Properties dialog box changes to display the details of the newly selected object. Note that if you make any edits, the window becomes modal, which means you must apply the changes (or cancel them) before you can select another object.

You can modify the properties of multiple classification rules, policing rules, and access rules. For other multiple selections, properties cannot be displayed.

Selecting Objects

Select an object by clicking it. The icon of a selected object appears shaded.

The Details pane only displays object details if you double-click an object.

You can also select multiple objects by:

- Selecting objects one by one by pressing **Control** and clicking each object.
- Where objects are listed in the Details pane and the Palette, you can select a range of objects by pressing **Shift** and clicking the first and last items in the range.

Note that you can use the multi-select technique to:

- Create a new map or network view that includes the selected network objects. See "Creating Subsidiary Networks and Maps" for more information.
- Select a number of network objects (devices, sites, segments, and so on) and include them in a subsidiary network.
- Set the properties for a number of classification, policing, or access rules. For more information, see *IP Service Activator QoS User's Guide*.
- Edit the common properties of a number of classification rules, policing rules, and access rules.
- Copy and paste multiple rules.
- Delete multiple objects.

Creating New Objects

You create new objects (where it is permitted) through an existing object's context menu or the context menu associated with a folder.

An object's context menu provides options to create new objects of a type that the object model allows for that object.

In addition, you can generally add new objects from a folder's context menu. For example, you can add new device types using the **Device Type** folder's context menu. The exceptions to this rule are the **Devices** and **Segments** folders on the Topology tab and the **System Hosts** folder on the System tab.

Any new object is linked as a child of the object from which it was created. It may also be automatically linked to other objects. For example, if you create a new user in a user group, the user is a child of the user group object and also the system object.

You can view an object's parentage in the Ancestry pane.

To set up the objects representing your network topology (devices, segments, interfaces, and so on), use the device discovery process, which discovers details of the network, represents it graphically, and updates the object model.

To create an object:

1. Select the relevant folder or object in the Hierarchy, Ancestry, or Details pane and select an **Add** command from the context menu.

   For example, to create a new classification, select the **Classifications** folder on the Policy tab, and then select **Add Classification** from the context menu.

   IP Service Activator opens the Properties dialog box for that object.

2. Enter the details in the dialog box.

   If you need help to complete a particular property page, press F1 or click the **Help** button on the dialog box.

Viewing and Editing Object Properties

You view and edit the attributes or properties that are associated with an object in the object's property dialog box. Depending on the object's type, the dialog box may contain any number of property pages, which are accessed by clicking the appropriate entry in the navigation tree on the left side of the dialog box.

Properties dialog boxes are modeless windows, which means that you can leave them open while you select other objects from the client. As you select objects, the Properties dialog box changes to display the selected object's details. However, if you edit any of the details displayed in the dialog, the window becomes modal and you must apply the changes before you can select another object. You can only have one Properties dialog box open at a time.

Every Properties dialog box includes the following command buttons:

- **OK**: Validates the input, applies changes that have been made, and closes the dialog box.
- **Cancel**: Closes the dialog box without applying the changes. This button is displayed only after changes are made in the dialog box but have not yet been saved by clicking **Apply**.
- **Close**: Closes the dialog box. This button is displayed only if changes made in the dialog box have been saved by clicking **Apply**.
- **Apply**: Validates the input, applies changes that have been made but leaves the dialog box open. This button is grayed out initially and becomes active only when changes are made.
- **Help**: Displays help relevant to the selected property page.

Some property pages include an edit toolbar, shown in Figure 2–4.

Figure 2–4 The Edit Toolbar

From left to right, the edit toolbar buttons include:

- **Set**: applies a change to the selected entry.
- **Add**: adds a new entry.
- **Delete**: removes the selected entry.
- **Up**: moves the selected entry up one position.
- **Down**: moves the selected entry down one position.

To display an object's properties, on the **View** menu, select **Properties**.

If the object has no children, you can also double-click the object in the Hierarchy or Details pane to open the Properties dialog box.

Resizing an Object's Properties Dialog Box

You can resize the Properties dialog box for an object by clicking and dragging the resize handle in the lower right-hand corner of the dialog box. This is useful if a dialog box contains information that is not immediately viewable.

Changes to the default dialog box size persist, that is, IP Service Activator retains your preference.

Data Validation in a Properties Dialog Box

Data input is validated when you click **OK** or **Apply**, or when you select a different property page.

Additional validation is performed on the entire database whenever you save or commit a transaction or when you run the **Validate** command from the **Tools** menu.

Viewing an Object's Configuration

Objects, such as networks, customers, VPNs, sites, and so on, are referred to as 'policy targets.'

These objects can have policy types applied. For information about policy types, see *IP Service Activator QoS User's Guide* or IP *Service Activator VPN User's Guide.*

The policy elements associated with an object can be listed in the Details pane by double-clicking on the object and selecting the relevant tab.

Tabbed options at the bottom of the pane allow you to select the policy details you want to view.

To view the configuration associated with an object, double-click the object in the Hierarchy, Ancestry, or Details pane.

Depending on the object type, the configuration policies and services that are associated with the object are listed in the Details pane. You can view different elements by selecting from the tabs at the bottom of the Details pane.

Linking Objects

You set up associations between objects by linking them.

IP Service Activator shows the link between two objects as a link object, which appears in the Hierarchy pane as a copy of the object you have linked.

Examples of valid links include:

- Linking a device to the proxy agent that manages it.
- Linking a device to a site.
- Creating and modifying a structure of accounts, by linking accounts to account groups.
- Creating and modifying a structure of traffic types, by linking traffic types to traffic type groups.
- Linking sites to a VPN object to create a VPN.

Unlike a standard object, a link object's context menu includes an **Unlink** option.

This section describes how to link objects using the drag and drop and copy and paste link techniques. You can drag and drop objects to create links between them working in the Hierarchy and Details panes, or by opening new domain management windows and dragging between windows.

The copy and paste link technique uses options from the **Edit** menu. This involves copying the object and selecting the **Paste Link** option over the object to which you want to link the first object.

You cannot link all object types; valid links are dictated by IP Service Activator's object model. If you are using drag and drop linking, an invalid link is shown by the mouse pointer, which changes over an invalid link target. If you are using **Copy** and **Paste Link**, the **Edit** menu's **Paste Link** option is not enabled.

Note that there are other indirect methods of linking objects. You can create links between objects when you select options on an object's property page. For example, when you select a class of service on a PHB group's property page, you effectively create a link between the PHB group and the class of service. IP Service Activator also creates some links automatically; for example, when you drag discovered devices on to a topology map.

If you have created a link between object A and object B, object A cannot be deleted. You first have to delete all instances of the link object.

Link objects look just like real objects in the Hierarchy pane. To check if an object is the real object or a link, check the context menu. If there is an **Unlink** option, the object is a link.

To tell if an object has any links to it, check the context menu. If the **Delete** option is grayed out, the object has links. You can then look in the Ancestry pane to see what parent objects exist - the object will have more than one parent.

To create a link using drag and drop:

1. Organize the display so that the objects to be linked are both visible.

   You can simultaneously display objects in the Hierarchy and Details panes or by opening and arranging two domain management windows.

2. Drag the source object and drop it on the destination object. The mouse pointer changes as you drag the item:

   - The pointer appears with a small arrow beneath it when you are over a valid destination to which the object can be linked.
   - The pointer appears as a circle with a bar through it when you are over an invalid destination to which the object cannot be linked.

   A new link object appears in the new position in the hierarchy tree.

To create a link using the **Copy** and **Paste Link** commands:

1. Select the source object and select **Copy** from the **Edit** menu, or click the **Copy** toolbar button.

2. Select the object to which you want to link the copied object and select **Paste Link** from the **Edit** menu.

   If the destination is not valid, the **Paste Link** option is not available.

   When a link is successfully created, a new link object appears in the new position in the hierarchy.

Unlinking Objects: Deleting Links

An object may be linked to other objects at multiple points in the IP Service Activator object model. Where an object is linked to another object, IP Service Activator displays the linked object as a child of the object to which it is linked. Therefore, an object may appear in multiple locations within the IP Service Activator client, as with the Site Nimes in Figure 2–5.

Figure 2–5 Objects Linked at Multiple Points in the Hierarchy Pane

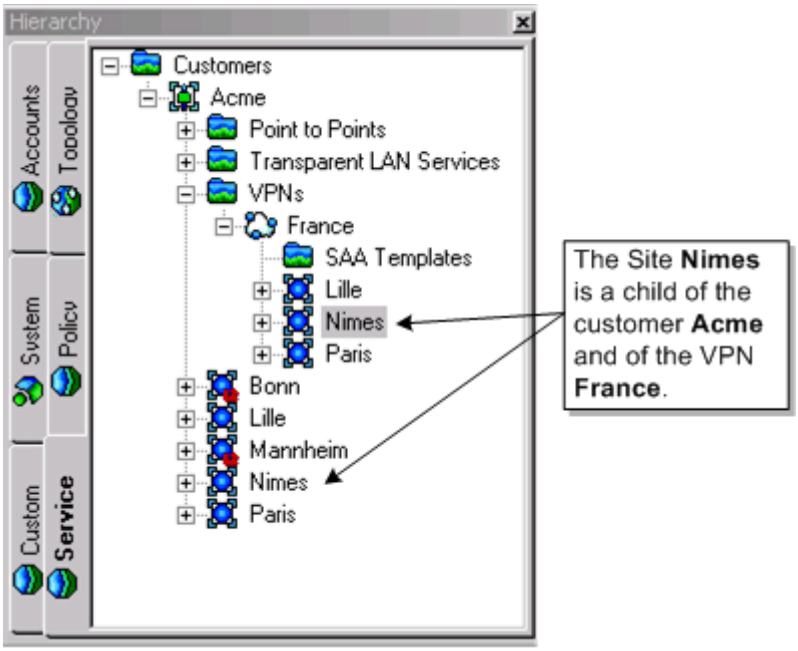

You can remove links between objects, for example if you have linked two objects by mistake or if you are making changes to the data.

Note: If an object has been linked to one or more parent objects, you must remove those links before you can delete the object. For example, if you have created a host account named 'zeus' and linked it to an account group, you cannot delete the zeus account until you have removed its link to the account group.

Note: If the **Unlink** option does not appear in the object's context menu, the object is a real object rather than a link object.

To check whether an object is a child of a parent object:

- Right-click the object to display the object's context menu.

  If the **Unlink** option is listed, the object is linked.

To unlink an object:

1. In a domain management window's Hierarchy pane, locate the parent object and select the link object listed below it.

For example, to delete the link between a proxy agent and a device, locate the proxy agent in the hierarchy tree and select the device link object, listed beneath the proxy agent.

2. From the **Edit** menu, select **Unlink** to unlink the object.

   The link is deleted and the link object is deleted from the hierarchy.

> **Note:** You can also drag the object onto a new object, which removes the existing link and links the object to a new parent.

Deleting Objects

You can delete objects from the object model, for example a router that does not physically exist in your network, or a rule component that is no longer used.

Note that you cannot delete system-defined (top-level) folders. These provide a means of grouping objects but are not objects themselves. See "Organizing Objects into User-defined Folders" to delete a user-defined folder.

> **Note:** Take care when deleting objects. IP Service Activator does not always display a confirmation message prior to deletion. If you delete something by mistake you can use the **Edit** menu's **Undo** command to restore it provided that you have not saved the current transaction since the object was deleted. You may also be able to rollback the transaction that performed the deletion.
>
> For more information about undoing commands during a transaction or about rolling back a transaction, see *IP Service Activator Concepts*.
>
> You cannot delete any object currently in use, for example, a Date and Time template that is assigned to a policy rule.

If you have dragged one object on to another, creating a link, the first object cannot be deleted until all links are deleted.

Remember that if you delete an object (such as a device or interface) from the map, you are actually deleting the object from the database.

To delete an object, select the object in the Hierarchy, Ancestry or Details pane, and, from the **Edit** menu, select **Delete**.

The object is deleted from the object hierarchy.

Organizing Objects into User-defined Folders

Most types of objects can be organized into user-defined folders.

- Customers: customer folders are created under **Customer** objects in the **Service** tab
- Sites: customer sites are created under **Site** objects in the **Service** tab
- Driver Scripts: driver script folders are created under the **Driver Scripts** folder in the **Custom** tab
- PHB Groups: PHBs and MQC PHBs can be organized into folders and sub-folders under the **PHB Groups** folder in the **Policy** tab.

- Classifications: classifications and classification groups can be organized into folders and sub-folders under the **Classifications** folder in the **Policy** tab.
- Classes of service: CoS can be organized into folders and sub-folders under the **Classes of Service** folder in the **Policy** tab.
- Roles: roles can be organized into folders and sub-folders under the **Roles** folder in the **Policy** tab.
- Traffic Groups: effectively folders organized under the **Traffic Types** folder in the **Policy** tab.

User-defined folders can contain sub-folders. Drag and drop operations can be used to move folders or their contents. You can also define folder permissions to restrict access to a subset of users for these folders.

To create a user-defined folder:

1. Select the correct tab, depending upon the user.
   - For customers, select the **Service** tab on a **Domain** detail window or **Network** detail window.
   - For sites, select the **Service** tab on a **Domain** detail window.
   - For driver scripts, select the **Custom** tab on a **Domain** detail window, **Network** detail window, or **Global View** window.
   - For PHB groups, classifications, classes of service, roles, and traffic groups, select the **Policy** tab on a **Domain** detail window.
2. Right-click on the parent object for the new folder. This can be the main system-supplied container folder (such as Customers, Driver Scripts or PHB Groups) or an existing user-defined folder.
3. From the context menu, select **Add Folder**. (For Traffic Types, select **Add Traffic Group**.)

   One of the following Folder dialog boxes opens, as appropriate:
   - Customer Folder
   - Site Folder
   - Driver Script Folder
   - PHB Group Folder
   - Classification Folder
   - Classes of Service Folder
   - Roles Folder
   - Traffic Group
4. Enter an identifying **Name** for the folder and, optionally, accompanying **Remarks**.
5. If you wish to restrict access to the folder, select its **Ownership** property page and specify permissions.
6. Click **OK** and commit the transaction.

To add an object to a user-defined folders:

1. Drag the object to the target folder. Objects can be included in only one folder.

To move an object out of a user-defined folder:

1. Drag the object to the new folder, or to the root folder for the objects.

To move a user-defined folder:

1. Click the folder and drag it to its new parent. The parent can be the root folder for the objects or another user-defined folder of the same object type. Contents (objects and subfolders) of the folder move with the folder.

To delete a user-defined folder:

1. Click the folder and, from the **Edit** menu, click **Delete**. When a user-defined folder is deleted, all of its contents (objects and subfolders) are deleted with it.

To rename a user-defined folder:

1. Right-click the folder and select **Properties** from the context menu.
2. Enter the new name for the folder in the **Name** field, and click **OK**.
3. Commit the transaction.

Navigation

IP Service Activator offers a range of techniques for navigating the client.

- View different object groups by selecting tabs in the domain management window.
- Navigate the hierarchy tree by expanding and contracting objects in the Hierarchy or Ancestry panes.
- Single-click on objects in the hierarchy tree to browse the hierarchy while keeping the view in the details pane static.
- View more information about an object in the Details pane by double-clicking on the object in the Hierarchy pane.
- Create different views of the policy domain simultaneously by opening and arranging new windows. See "IP Service Activator Windows" for more information about working with windows.
- To open a duplicate of the current explorer window, choose **New Window** from the **Windows** menu.
- To open a new explorer window which has the currently selected object at the top of the hierarchy, choose **Open** from the **View** menu or from the context menu. The selected object becomes the title of the new window.

About the Toolbar

Figure 2–6 shows the toolbar at the top of the main window that includes icons that allow you to access the most frequently used commands.

Figure 2–6 The Toolbar

Table 2–2 describes the function of each toolbar icon as they appear, from left to right.

Table 2–2 Toolbar Icons

Icon	Description
Validate	Checks all the data in the object model and reports any errors in the Current Faults Pane.
Save	Saves the current transaction to the database. The changes held in the transaction are not applied until the transaction is merged with the local object model and then committed.
Commit	Saves the current transaction to the database and propagates any associated device configuration to the network.
Schedule	Displays the Transaction Properties dialog box, allowing you to store the current transaction for implementing at a set date and time.
Abort	Aborts the current transaction, discarding all changes.
Copy	Places a copy of the selected object on the Windows clipboard. This has the same effect as the **Copy** command in the **Edit** menu. To copy text from property pages or dialog boxes, use the **Copy** command from the context menu.
Paste	Pastes the object in the clipboard to the selected point. Can be used to copy rules. To paste text from property pages or dialog boxes, use the **Copy** command from the context menu.
Undo	Reverses the previous action. You can click on this button more than once to reverse actions within the current transaction.
Redo	Performs an action again, if you have selected Undo and you then decide you didn't want to undo the action. You can click on this button more than once in order to reverse Undo operations within the current transaction.
Delete	Deletes the currently selected object.
Unlink	Removes the link between the currently selected object in the object model and its parent.
Properties	Displays the Properties dialog box for the selected item, allowing you to view or edit the information. Opens a new explorer window, with the currently selected item as the topmost item in the hierarchy
Synchronize panes	Synchronizes the display in the Hierarchy pane with the Details pane. The parent of the item selected in the Details pane is highlighted in the Hierarchy pane.
Back	Returns to the previous display from a previously-navigated browse sequence.
Forward	Returns to the next display in a previously-navigated browse sequence
Configuration	Displays details of the configuration that applies to the currently selected object, including rules, PHB groups, VPNs, and driver scripts.
Report view	Displays the contents of the Details pane in the form of a list of objects, together with additional information about the object.
Compact list view	Displays the contents of the Details pane in the form of a list of objects.
Map view	Displays the network map in the Details pane.
Status context	The map display uses different colors to show the status of devices and interfaces.
Policy context	The map display uses different colors to show the status of devices and interfaces.

Table 2–2 (Cont.) Toolbar Icons

Icon	Description
Print	Displays a standard Windows print dialog box, allowing the selected network map to be printed. This is only available when the map is displayed.
Help	Displays the contents page of the on-line help.

Synchronizing Panes

When you double-click on an object in the Hierarchy pane, the Details pane synchronizes automatically and the selected object's details are displayed in the Details pane. However, you can also click on objects to browse the hierarchy tree but leave the view in the Details pane static.

The Synchronize Panes feature enables you to resynchronize the view shown in the Hierarchy pane with the Details pane. It returns you to the tab whose details are shown in the Details pane, with the tab in the same state as last viewed.

To synchronize the Hierarchy pane with the Details pane:

- Click the **Synchronize Panes** icon on the toolbar.

IP Service Activator synchronizes the Hierarchy pane with the Details pane.

Note: Synchronization only occurs if you have navigated objects in the Hierarchy pane using the single-click technique.

Using the Navigation Icons

Whenever you open a new window and browse through the object model, IP Service Activator keeps a record of the object views you have visited. You can move through the recorded browse sequence using the **Up**, **Back**, and **Forward** icons.

The **Up** icon is only enabled when you drill down through a map on the **Topology** tab. When selected, it takes you up one level in the topology hierarchy and displays details of the selected object's parent. However, if you have a segment selected, the **Up** icon is not enabled. This is because a segment may have a number of connected interfaces and therefore has multiple parents.

To navigate a recorded browse sequence:

- Use the **Back** and **Forward** icons to step backwards and forwards through the browse sequence.
- Use the **Up** icon to return to a previous map view after you have drilled down a level.

Changing Views

The Details pane can display details of selected objects in several formats. You can switch between views using toolbar icons:

- **Report View**: displays the contents of the Details pane in the form of a list of objects, together with additional information about the object's properties. You can sort on any column by clicking on its title.
- **Compact List View**: displays the contents of the Details pane in the form of a multi-column list of objects.

- **Map View**: displays the network topology map or the VPN map in the Details pane. This button is only available when you have selected an object that has an associated map, that is, a network object or a VPN object.
- **Status Context**: displays the map view showing the status of all network elements.
- **Policy Context**: displays the map view showing the policy role of all network elements.

You can also change the type of information that is displayed on the topology map, viewing information about the status of the map's devices or the policy implemented at each device. See "How Objects are Represented" for more information.

Changing the Appearance of the Client

IP Service Activator gives you flexibility over the appearance of the client. You can hide and show window panes, arrange panes or move elements around the window. The only window element that you cannot hide is the Details pane.

Hiding and Showing Windows and Window Elements

You can hide and show windows and window elements by selecting options from the **View** menu. You can hide or show the following elements:

- Toolbar
- Status bar
- Hierarchy tree pane
- Ancestry tree pane
- Map palette
- Global Setup window
- Current faults pane
- Statistics summary pane
- System statistics window

To hide or show a window element, from the **View** menu, select the element that you want to hide or show.

An element is displayed if its name is checked.

Changing the Size and Position of a Pane

You can change the size of any of the displayed panes and, for selected panes, drag them to a new position.

To change the relative size of panes:

- Click on and drag the vertical and horizontal bars separating the relevant panes.

To move and dock a pane, do one of the following:

- On the Hierarchy, Ancestry or Palette pane, click on the pane's title bar and drag it to a new position.
- On the Current Faults pane, click on the pane's docking bar and drag the pane to a new position.

Depending on the pane's new location, IP Service Activator may change the pane's orientation and size.

If you drag the pane outside the domain management window, it becomes a separate, floating window. To return the pane to the window, position the pane over the window's edge.

Changing the Hierarchy Pane's Tab Position and Style

You can customize the appearance of the Hierarchy pane's tabs by selecting their location on the pane. If tabs are displayed at the top or bottom of the pane, you can also specify whether they appear in a single line or multiple lines. The default is multiple lines.

To change tab position and style:

1. Right-click on the Hierarchy pane's title bar and select **Display Tabs** to display the pane's context menu.
2. Select **Top**, **Bottom**, **Left**, or **Right**.

   The tabs move to the position you selected.
3. Right-click on the Hierarchy pane's title bar and select **Multilined**.

   The **Multilined** option is only displayed if tabs are displayed at the top or the bottom of the pane.

   The tab style changes depending on whether the **Multilined** option is selected or deselected.

Searching IP Service Activator

IP Service Activator provides the following options for searching:

- Search for objects based on a text string

  The text string may be in the form of a regular expression. You can choose whether to expand or restrict the search to policy objects, including inherited policy objects. You can also restrict the search by object type, state and/or role.
- Search for concrete configuration objects

  This search locates concrete rules, PHB groups, driver scripts, VPNs, and CCCs. You can restrict the search by state.

Conducting a Text-based Search

A text-based search is based on a string that may include a regular expression. The search operates from a user-selectable point in the hierarchy. IP Service Activator attempts to match the search string against all attributes of an object. This means, for example, that you can search for a device by name, role or IP address, or a rule by its name, status or associated traffic type.

The search text may include any alphanumeric or punctuation characters up to a maximum of 64 characters.

By default, IP Service Activator attempts to match the search text at any point in a text string. For example, the search text '165' will be matched against 168.165.0.4. Alternatively, you can specify that the search text is matched against whole words only. You can specify whether the search is case sensitive.

IP Service Activator remembers previous search start points and displays them in a pull-down list in the Find dialog box. This list is cleared if the client is stopped and restarted.

Note the following:

- IP Service Activator searches the contents of the Object Model. Note that the System Logs are not held in the Object Model, and you therefore cannot search their content from the client.
- You can search within a single domain only.

> **Note:** For complete dialog box and property page descriptions, see IP Service Activator online Help.

To conduct a text-based search:

1. From the **Tools** menu, select **Find**.

   The **Find** dialog box opens.

2. In the **Find What** field, specify the search string.
3. Optionally, select options including **Search From**, **Match Whole Word Only**, **Match Case**, **Regular Expression**, **Match Configuration**, **Match Configuration Only**, and **Match Inherited Configuration**.
4. If you want to restrict the search by object type, state or role, select the **Advanced** tab and select from options, including **Match Object Type**, **Match Object State**, and **Match Object Role**.
5. Click **Find**.

   IP Service Activator searches for the specified search criteria. As matches are found, the search results are added to the list in the lower half of the Find dialog box.

> **Note:** Depending on the amount of information held by IP Service Activator, a search may take some time to complete. You can stop the search by clicking the **Cancel** button that is displayed while the search is in progress.

When the search is complete the results are shown in the lower half of the Find dialog box.

By default, found objects are sorted by object name. You can:

- Change the sort order criteria by single-clicking on a column title.
- Display a found object in the Hierarchy pane by double-clicking on it. IP Service Activator highlights the object in the Hierarchy pane and, where relevant, lists children of the found object in the Details pane.
- Open a new window on a found object by selecting **Open** from the object's context menu.
- View an object's properties by selecting **Properties** from the object's context menu.

For information about searching for concrete objects, see *IP Service Activator Concepts*.

Searching Committed Transactions

IP Service Activator supports the ability to search through committed transactions. See "Searching Committed Transactions" for details.

Printing

You can print the contents of the Details pane to produce, for example, a printout of the system log file or topology map. Each page has a ruled border with information about the page content printed round the border edge. Details of the page layout, including margin size, can be specified and the printed layout previewed.

> **Note:** When printing a topology map that occupies multiple pages, there is some overlap between pages.

For information about printing reports, see *IP Service Activator Network and SLA Monitoring Guide*.

To print the Details pane:

1. From the **File** menu, select **Print**.

   The Print dialog box opens.

2. Adjust the print settings if required.
3. Click **OK**.

To define the print page setup:

1. From the **File** menu, select **Page Setup**.

   The Page Setup dialog box opens.

2. Select the **Paper**, **Orientation**, and **Margins** that you want.
3. If you want to select a printer, click the **Printer** button, select the printer, and click **OK** to confirm the printer selection.
4. Click **OK**.

To display a print preview of the Details pane:

1. From the **File** menu, select **Print Preview**.

   IP Service Activator displays a preview of the pane printout. You can move between pages using the **Next Page** and **Prev Page** buttons. You can view the image in more or less detail using the **Zoom In** and **Zoom Out** buttons.

2. If you want to print the previewed image, click **Print** to display the Print dialog box and **OK** to confirm the print request.
3. If you want to close the previewed image, click **Close**.

3

Defining and Applying Roles

This chapter introduces the concept of roles in Oracle Communications IP Service Activator and describes their function in applying policy and measurement to the network.

The chapter:

- Provides an overview of roles
- Describes the role types that feature in IP Service Activator – system and user-defined roles
- Explains how to associate a role with a policy target

About Roles

A role is a means of grouping a set of policy targets that should 'attract' the same policy-based configuration, as shown in Figure 3–1.

Figure 3–1 A Role Attracting Policy Targets and Policy Elements

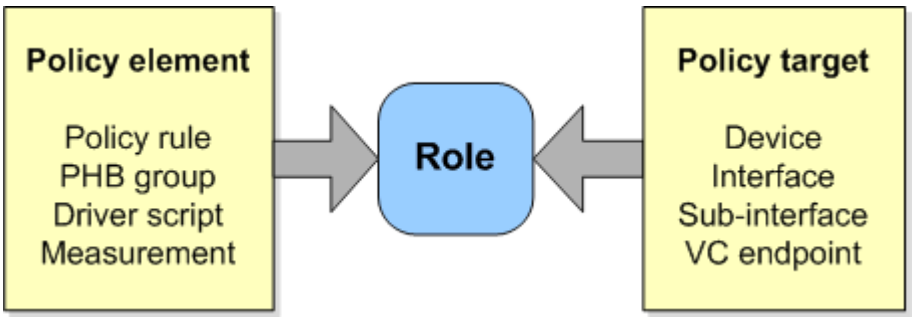

You can associate roles with:

- Some policy targets: a device, interface, sub-interface or VC endpoint
- Policy elements: a rule, PHB group, driver script or measurement policy element (measurement or collector parameters)

IP Service Activator applies QoS and measurement policy elements to policy targets only where their roles match.

This means that you can target QoS or measurement policy to specific points in the network. Figure 3–2 shows this concept using a policing rule.

Figure 3–2 Targeting Policy Elements Using a Policing Rule

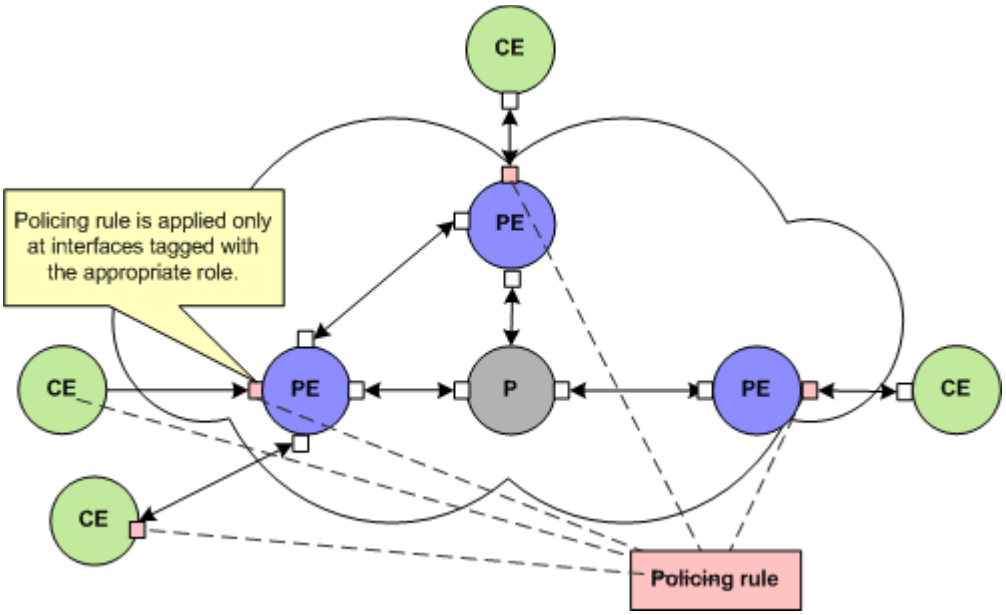

The ability to associate a role with any policy target down to the VC endpoint level provides a fine degree of control. For example, you can tag a group of sub-interfaces with a role to apply policy at the sub-interface level, independent of their parent interfaces.

IP Service Activator provides a set of system-defined roles that follow the DiffServ model but you can create additional user-defined roles. For example, it is possible to create roles that group interfaces according to their bandwidth and apply different policy to each set of interfaces in accordance with their bandwidth capacity.

Roles must be assigned to policy targets manually. You must assign a role manually for each device.

For information on associating roles with policy elements, see *IP Service Activator QoS User's Guide*. For information on associating roles with collector and measurement parameters, see *IP Service Activator Network and SLA Monitoring Guide*.

Both system and user-defined roles are further subdivided into device and interface roles:

- A device role can be assigned to devices.
- An interface role can be assigned to interfaces, sub-interfaces or VC endpoints.

System and User-defined Roles

IP Service Activator recognizes two types of role:

- System-defined: a set of device and interface roles that support the DiffServ policy model.
- User-defined: a role that has been created by a user.

Both system and user-defined roles are available across all domains.

Viewing the Roles Defined in IP Service Activator

You can list the device and interface roles that are currently defined in IP Service Activator on the **Policy** tab beneath the **Roles** folder. Figure 3–3 indicates which role each role icon represents.

Figure 3–3 Role Icons and Definitions

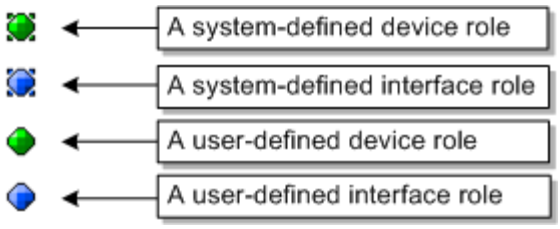

System-defined Roles

System-defined roles follow the DiffServ model. See *IP Service Activator Concepts* for details of the policy model.

You can assign one system-defined role and multiple user-defined roles to a policy target. Policy elements can be associated with only one system-defined role and one user-defined role.

> **Note:** If you are implementing MPLS VPNs or measurement-only VPNs with IP Service Activator you must use system-defined roles. For information about how roles are used within MPLS VPNs, see *IP Service Activator VPN User's Guide*.

The following system-defined roles are provided in IP Service Activator. The roles are divided into two types: Device and Interface.

- Device roles:
  - Access: Devices that provide access to the core network or WAN from an external subnet or customer LAN. Equivalent to the Customer Edge (CE) role in MPLS VPNs.
  - Gateway: Devices on the edge of the core network or WAN that directly connect to the local or customer access device. Equivalent to the Provider Edge (PE) role in MPLS VPNs.
  - Core: Devices within the core network. Equivalent to the Provider (P) role in MPLS VPNs.
  - Shadow: Devices dedicated to Service Assurance Agent (SAA) measurements in a Service Provider's Points of Presence (POP).
- Interface roles:
  - Local: An inbound interface from a customer site.
  - Access: An interface that connects Gateway and Access devices.
  - Core: An interface that connects two Core devices, or a Core and a Gateway devices.
  - Disabled: A non-managed interface.

The points in the network at which these roles apply in the DiffServ model are shown in Figure 3–4.

Figure 3–4 Network Points at Which the Interface Roles Apply

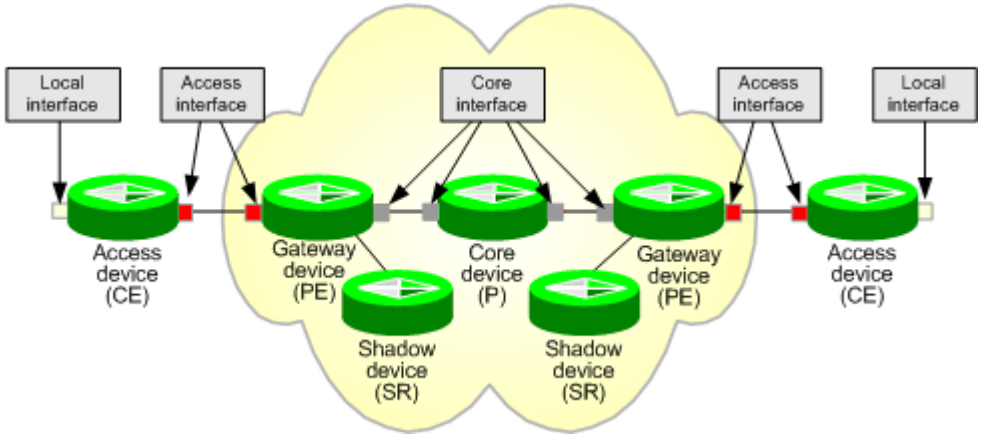

For information about the Shadow role and SAA measurement, see *IP Service Activator Network and SLA Monitoring Guide*.

Note: You cannot delete system-defined roles.

User-defined Roles

You can create user-defined device and interface roles and assign any number of user-defined roles to a policy target. By assigning multiple roles to a policy target, it becomes a member of several role groups.

You can assign user-defined roles to policy targets in combination with system-defined roles.

The number of roles you need to create depends on the network set-up and the number of variables that affect policy. For example, you may need to create roles that classify policy targets by link capacity, device function, customer and service package. For an example based on user-defined roles, see *IP Service Activator QoS User's Guide*.

Note: The ability to create roles depends on your user security level. See "User Access Levels and the Client" for more information.

To create a role:

1. On the **Policy** tab, select the **Roles** folder.
2. Select **Add Device Role** or **Add Interface Role** from the context menu.

   The Device or Interface Role dialog box appears.
3. Enter details including **Name** and **Remarks**.

Deleting User-defined Roles

You cannot delete a role that is linked to a policy element or to a policy target – unlink the role before deleting it. You cannot delete system-defined roles.

To delete a role:

1. On the **Policy** tab, select the **Roles** folder.
2. Select the device or interface role that you want to delete.
3. From the context menu for the role, select **Delete**.

Assigning a Role to a Policy Target

You must assign one or more roles to each device and interface, sub-interface or VC endpoint to be managed in order to define the points in the network at which services will be configured and policy applied.

You can assign one system-defined role to a device, interface, sub-interface or VC endpoint, and any number of user-defined roles.

You can apply a role by:

- Manually assigning the role to an object

When you apply a role manually, the role applies to the selected policy target only.

Assigning Roles

Assigning roles to the devices and interfaces in the domain is an essential part of setting up IP Service Activator.

You can apply a role to a policy target manually using the object's properties dialog box.

> **Caution:** If role assignment rules are subsequently applied, where a policy target matches the parameters specified by a role assignment rule, a manually-assigned role is overridden (and configuration associated with the previous role is removed). The exception to this is the system-defined interface role Disabled, which is never overridden. Role assignment rules are not recommended.

To assign and unassign a role manually

1. Open the **Topology** tab and select the device, interface, sub-interface, or VC endpoint to which you want to assign a role.
2. From the device or interface's context menu, select **Properties** and select the **Role** property page.

   The property page lists the currently-defined device or interface roles.
3. If you want to assign a role, select the check box associated with the role name.

> **Note:** You can only assign one system-defined role. Selecting another pre-defined role to the list replaces the previously-assigned role. You can assign any number of user-defined roles.

4. If you want to unassign a role, deselect the check box associated with the role name.

4

Setting Up Domain Information

This chapter explains the tasks you need to do immediately after installation of Oracle Communications IP Service Activator in order to set up domains.

The chapter describes how to:

- Create domains and set up proxy agent assignment for the devices within the domain
- Load files containing basic configuration data
- Open domains

Setting Up Domains

Oracle Communications IP Service Activator uses a concept of domains to define the logical networks to be managed. You can create multiple domains to represent managed networks: one domain for each AS region to be managed. No domains exist when you first install IP Service Activator, so the first step is to create one or more domains.

To create a new domain:

1. Select **New Domain** from the **File** menu. The Domain dialog box opens.
2. Enter an identifying name for the domain.
3. Specify the type of domain:
   - **Public** if the network only uses public IP addresses.
   - **Private** if you plan to manage independent networks with overlapping address space.
   - **MPLS VPN** if you are going to set up MPLS-based virtual private networks.
4. For MPLS VPN domains, specify whether interfaces on Provider Edge routers that connect to the Customer Edge routers use public or private addresses.

   By default, they are assumed to use unique, public IP addresses. If they use private IP addresses, clear the **Public PE to CE Addresses** check box.

   For further configuration, see the cartridge or driver guide that applies to your configuration.
5. In the **Manual Config** field, specify the default action within this domain for detecting manual configuration. The **Manual Config** field setting does not apply to the Network Processor. The system only checks to see if changes have been made to the commands that may be configured by IP Service Activator, as these may affect the operation of the system:

- **Delete**: If manual configuration is detected, the device drivers will issue commands to reconfigure the device. No warning is output regarding the detected manual configuration.

Caution: Oracle strongly recommends that you use only the **Delete** setting.

- **Warn and delete**: If manual configuration is detected, a warning message (Message 3203) is output and the device drivers will issue commands to reconfigure the device.
- **Fail and don't delete**: If manual configuration is detected, a critical fault (Message 3494) is raised. The device status is set to Intervention Required but is not reconfigured.

Note: The domain-level manual configuration settings can be overridden for specific devices (see "Setting Manual Configuration Detection").

IP Service Activator never deletes manually pre-configured VRF tables even if you select the **Delete** or **Warn and delete** options.

---

6. If you want to own the domain and set permissions on it for yourself, members of your user group and other users, select the **Ownership** property page.

Note: If you are setting up VPNs you also need to set parameters on the **VPN BGP**, **VPN MPLS** and **ASN** property pages. For full details of how to set up VPNs, see *IP Service Activator VPN User's Guide*.

Manual Configuration with Manual Config Field set to Delete

On the Domain properties page of the Domain dialog box, the **Manual Config** menu lets you select how IP Service Activator deals with manually applied configuration statements that it encounters on the devices. The **Manual Config** field does not apply any configurations to the Network Processor. See "Setting Up Domains"for more information.

Caution: Oracle recommends that you use only the **Delete** option.

Device Driver Handling of Manually Configured Objects when Delete Selected

This section discusses how IP Service Activator deals with manually applied configuration statements on devices that are managed by a Device Driver. See "Network Processor Handling of Manually Configured Objects when Delete Selected" for information about how this process relates to network processor or cartridge managed devices.

If the name of a manually configured object on a managed device exactly matches a name that already exists in the IP Service Activator object model, IP Service Activator assumes control of the configuration of that object. The parameters stored in the object model are used to configure the object on the device so that it aligns with the object model. This overwrites the manually applied configuration of that object if it differs from the object model.

If the name of a manually configured object exactly fits the default naming convention used by IP Service Activator for such objects, IP Service Activator assumes that the object was deleted from the object model but erroneously left on the device. IP Service Activator then deletes the manually configured object from the device, so that the device configuration matches the object model.

If a manually configured object's name does not match any name in the object model, and does not fit the default naming convention used by IP Service Activator, the object is left on the device and its configuration is not altered.

Network Processor Handling of Manually Configured Objects when Delete Selected

This section discusses how IP Service Activator deals with manually applied configuration statements on devices that are managed through cartridges and the Network Processor.

IP Service Activator does not maintain an internal model of the state of devices managed through the Network Processor. Configuration for particular objects are applied to the device and if there is a manually configured object already on the device with the same name, it is overwritten, or potentially augmented, depending on the type of object and the particular configuration statements involved.

If the name of a manually configured object on the device does not match any of the names in the IP Service Activator Configuration, the object is left on the device and its configuration is not altered.

Setting the Default Loopback ID Value for Discovery

On the Domain dialog box, **VPN BGP** property page, you can specify a value for the **Loopback ID** field. The Loopback ID value is used to create a loopback interface name by appending it to the name 'loopback'. For example, if the Loopback ID is 0, the loopback interface name created is 'loopback0'. When a device in this domain is discovered, a check is made to see if a loopback interface matching this text string exists. If it does, the IP address of the loopback interface is stored with the device information. The Loopback ID value can be overridden on a per-device basis (on the Device dialog box) and on a per-discovery basis (on the Topology dialog box).

Setting Up Proxy Agent Assignment

All devices in the domain that are to be managed by IP Service Activator must be assigned to a proxy agent. It is the proxy agent that controls when and what type of configuration is to be applied to a specific interface.

Although it is possible to assign devices to proxy agents manually, it is generally performed automatically during device discovery.

Assigning a Proxy Agent to a Domain

If you have a distributed installation with multiple proxy agents and you are creating multiple domains, you can assign specific proxy agents to each domain.

A device can only be assigned to a proxy agent that is either global or has been assigned to the domain that the device is in.

To assign a proxy agent to a domain:

1. Select a domain from the global setup window and select **Properties** from the context menu.

2. Select the **Proxy Agent** property page. The list shows all the proxy agents currently installed.
3. Select one or more proxy agents to be used within the domain by clicking the check box associated with the proxy agent.

Repeat these steps to assign specific proxy agents to specific domains.

> **Note:** If you do not explicitly define proxy agents for each domain, all proxy agents remain global and devices within any domain can be assigned to them.

Setting Up Proxy Agents for Automatic Assignment

Devices can be assigned to a proxy agent automatically whenever devices are discovered or rediscovered. You can configure IP Service Activator either to assign all devices to one proxy agent or to assign the devices equally to a number of proxy agents.

Devices are assigned only to proxy agents that are:

- Specifically assigned to the domain that the device is in, or are global (that is, not assigned to any domain) and
- Defined as active for auto-assignment

To check that a proxy agent is active, display the proxy agent properties by selecting the relevant proxy agent on the **System** tab in the **System Hosts** folder and choosing **Properties** from the context menu. Ensure the **Auto device assignment** is set to **On** (this is the default setting).

To configure automatic proxy agent assignment within a domain:

1. From the **Tools** menu, select **Options**.

   The Options dialog box opens.
2. Under **Auto proxy assignment**, select either **First** or **Load balance**:
   - **First** assigns each new device discovered to the first active proxy agent. This is the default setting.
   - **Load balance** allocates devices to multiple proxy agents with an equal allocation to all active proxy agents in order to balance the processing load between them.
   - If **Off** is selected, devices are not assigned to proxy agents automatically. Before they can be managed you will need to link them manually.

Only supported devices, that is, those that are included in the **DeviceTypes.cfg** file in the **Config** directory, can be automatically assigned.

Working in multi-proxy environments

When running in a multi-proxy environment:

- If you discover a device and fetch capabilities under a Domain which has both a Network Processor and a Proxy Agent proxy active, you will not be able to tell which fetched the capabilities. In this case, assign the device to the desired proxy, and then reset and re-fetch the capabilities.

Loading Policy Configuration Data

The next step is to install one or more set-up files that create standard policy configuration data in the IP Service Activator database – basic policy components and, if required, sample rules and PHB groups.

> **Note:** Oracle recommends loading the default configuration policy data. If you do not load the data, you have to create all the basic component data manually. Loading configuration policy data is mandatory if you want to use configuration policy.

The following configuration files are supplied in the **SamplePolicy** directory:

- **default.policy** creates some basic policy data, including:
  - Packet markings representing IP Precedence values 0-7
  - Gold, Silver and Bronze classes of service
  - Packet marking based traffic types representing the Gold, Silver and Bronze classes of service and port-based traffic types representing the most common TCP and UDP port numbers
  - Classifications based on the traffic types
- **advanced.policy** creates additional policy data:
  - Packet markings based on the full range of DiffServ codepoints and MPLS experimental bits
  - Packet marking based traffic types representing DiffServ codepoints and MPLS experimental bits
  - Port-based traffic types representing the most common IP protocols
  - Classifications based on the packet marking traffic types

  This data is useful if your routers support the full range of DiffServ codepoints and/or MPLS experimental bits.
- **Rule_and_PHB.policy** includes some example policy rules and standard PHB groups. If you wish, you can base your own rules and PHB groups on these examples.

The files must be loaded in the order in which they are listed above.

The following files may also be loaded:

- **ConfigletPolicyTypes.policy** loads the configlets, which are needed by CTM.
- **ConfigurationManagementPolicyTypes.policy** loads other config policies needed by Configuration Manager.
- **ExtensionPackAPolicyTypes.policy** contains standard definitions for:
  - Extension Pack A Configuration Policy Types
  - ATM PVC Vc Class
  - Save Running Config
  - Static NAT
- **ExtensionPackBPolicyTypes.policy** contains standard definitions for:
  - Extension Pack B Configuration Policy Types

  - VRF Custom Naming
  - VRF Export Route Filter
- **ExtensionPackCPolicyTypes.policy** contains standard definitions for:
  - Extension Pack C Configuration Policy Types
  - ATM Vc Class
  - BGP CE
  - Extended ACL
- **GeneralPolicyTypes.policy** contains standard definitions for:
  - General Configuration Policy Types
  - Banners
  - IP Pools
  - Key Chain
  - Prefix List
  - SNMP Community
  - SNMP Host
  - Static Route
  - User Authentication
  - User Data
  - Traffic Statistics
- **InterfaceManagementPolicyTypes.policy** contains standard definitions for:
  - Interface Creation Configuration Policy Types
  - Interface Decoration Configuration Policy Types
  - Subinterface Creation Configuration Policy Types
  - Channelized Interface Creation Configuration Policy Types
  - Other Interface Management Configuration Policy Types
- **IPMulticastPolicyTypes.policy** contains standard definitions for:
  - IP Multicast Configuration Policy Types
  - Mulitcast Interface types can be installed by loading the
  - MulticastInterfacePolicyTypes.policy file
- **IPSecPolicyTypes.policy** contains standard definitions for:
  - IPSec Configuration Policy Types
  - Customer IPSec
  - Public IPSec
  - IPSec
- **JuniperJUNOSQoSPolicyTypes.policy** contains standard definitions for:
  - Juniper JUNOS QoS Configuration Policy Types
  - Juniper JUNOS COS Attachment

- **juniper.policy** defines MPLS packet markings, classes of service and an example PHB group for configuring WRR on Juniper M-series devices. For more information, see the *IP Service Activator Juniper M-series Driver Support Guide*.
- **LSPPolicyTypes.policy** contains standard definitions for:
  - LSP Configuration Policy Types
  - LSP Tunnel
- **MulticastInterfacePolicyTypes.policy** contains standard definitions for Multicast Interface Configuration Policy Types to be used with VPN and IP Multicast Policy Type services.
- **Role_Assignment_Rules.policy** defines a set of role assignment rules that allocate system-defined roles to devices and interfaces. Do not use role assignment rules - assign roles manually.
- **RoutePolicyPolicyTypes.policy** contains standard definitions for:
  - Route Policy Configuration Policy Types
  - BGP Route Policy
  - VRF Route Policy
- **ServiceAssurancePolicyTypes.policy** contains standard definitions for:
  - Service Assurance Configuration Policy Types
  - Collector Parameters
  - Netflow Parameters
  - RTR Responder
- **VLANPolicyTypes.policy** contains standard definitions for:
  - VLAN Configuration Policy Types
  - MGMT VLAN Interface
  - VLAN Interface
  - VLAN Definition
- **VPNMulticastPolicyTypes.policy** contains standard definitions for:
  - VPN Multicast Configuration Policy Types
  - Mulitcast Interface types can be installed by loading the
  - MulticastInterfacePolicyTypes.policy file

The **SharedPolicyData.policy** file is loaded automatically at system start-up. It defines a set of commonly-used IP protocols which are available in any domain you create. You only need to load this file if the IP protocols are deleted or edited incorrectly.

All of the above files are domain-specific – that is, they must be loaded into each domain in which you wish to use their information.

To load a policy configuration file:

1. Select a domain from the global setup window and select **Properties** from the context menu.
2. On the Domain dialog box, select the **Setup** property page.
3. Click **Browse** to view the available configuration files in the SamplePolicy folder.

4. Select the file to load and click **Open**. A brief explanation of the file appears in the File Information box.
5. Click **Load** to load the selected file and create the data. Note that files must be loaded in the following order:
   a. **default.policy**
   b. **advanced.policy**
   c. **Rule_and_PHB.policy**

> **Note:** It is important to load the files in order. Always load **default.policy** first, and do not load **Rule_and_PHB.policy** if you have already created PHB groups.

Opening the Domain

To open the domain on which you are going to work, double-click on the relevant domain on the **Domains** tab, or select the domain and select **Open** from the context menu.

A new domain management window opens.

5

Discovering and Setting Up the Network

Before you can apply services or set up policies, you need to set up an accurate representation of the network to be managed. Oracle Communications IP Service Activator can automatically discover all IP addressable network elements within your network, that is, devices, hosts and network segments, and set up appropriate information. After you discover the devices to be managed, you can set up maps showing the network topology.

This chapter includes the following:

- Introduction to network discovery
- Steps you need to take before running device discovery, including setting the correct domain information
- The initial discovery of the network, including different methods of network discovery
- Steps you need to take after device discovery, to ensure that the information is complete

Introduction to the Discovery Process

During network discovery, IP Service Activator finds out details of devices, hosts, and network segments within the domain that you are managing. It also retrieves interface capabilities of each device and, where possible, sets up the information that IP Service Activator needs in order to manage the devices. The discovery process includes the following stages:

- Discovering devices, segments, and hosts
- Assigning devices to proxy agents
- Assigning policy roles to devices and interfaces
- Setting up the security parameters that IP Service Activator needs to configure devices
- Discovering the capabilities of devices and interfaces: the VPN, QoS, security, and measurement options that are available

Discovering the Network

The topology discovery process finds out information about the network topology by interrogating network nodes using SNMP. Given an IP address or DNS name, IP Service Activator discovers details of the node. If the node is classified as a device (that is, it forwards IP packets; for example, a router or a Layer 2 switch) details of

connected network segments, interfaces and any PVC endpoints are obtained. From a segment, IP Service Activator can find further directly connected nodes, that is, hosts and devices.

The discovery process is controlled by a set of SNMP parameters; default values are assumed by IP Service Activator, but these can be overwritten if required.

The way in which discovery works depends on whether public or private IP addresses are used within the domain. See "Before Running Device Discovery" for more information.

Assigning Devices to Proxy Agents

All devices in the domain that are to be managed must be assigned to a proxy agent. The proxy agents are the IP Service Activator components that manage the low-level device driver commands for each device type.

Assigning devices to proxy agents is generally performed automatically during device discovery. You can set up IP Service Activator either to assign all devices to one proxy agent or to assign the devices equally to all active proxy agents. See "Setting Up Proxy Agent Assignment" for more information.

Assigning Roles to Devices and Interfaces

All devices and interfaces to be managed must be assigned a role in order to define the points in the network at which service configuration or policy will be applied. Devices can be classified with any system-defined and/or user-defined roles that have been set up within the domain.

Setting Up Device Security Parameters

For each device, security parameters, such as user IDs and passwords must be specified to allow IP Service Activator to configure the device. This information must be set up prior to discovery because IP Service Activator requires write access to routers before it can obtain the device and interface capabilities.

Setting up standard security information applies it automatically to all discovered devices. This can save time if you use the same access methods and passwords for a number of devices in your network. See "Defining Default Security Options for Discovery" for more details.

Discovery Capabilities

At the end of the discovery process, IP Service Activator attempts to discover the capabilities of each device and its interfaces. These capabilities dictate the VPN, QoS, security, and measurement options available. See "Checking Capabilities" for more information. If IP Service Activator is unable to discover capabilities, for example because security parameters are incorrectly set up, you can discover the device capabilities at a later point. See "Refetching Device Capabilities" for more information.

Note: The discovery process can open a large number of file descriptors on Solaris systems. The size of the fd_set used is FD_SETSIZE, which is hard-coded to be 1024 on Solaris (in /usr/include/sys/select.h). By default, IP Service Activator can use half of this maximum. An alternative value can be set by specifying the following command-line parameter to the policy server:

```
-SnmpMaxSessions <xxx>
```

where xxx is the maximum number of file descriptors, between 1 and 1024, with the default value being 512.

BGP Autonomous System Discovery

Each router managed by IP Service Activator has a BGP Autonomous System (AS) number. IP Service Activator supports a default AS at the domain level but individual routers within this domain may have a different AS than the domain default. A domain that contains devices belonging to different Autonomous Systems is referred to as a Multi-AS domain.

IP Service Activator supports Multi-AS domains by actively discovering the BGP AS number configured into each device as it is discovered, or re-discovered, in the system. All BGP provisioning on a device will use the last discovered BGP AS number. If a device changes its Autonomous System membership, it must be re-discovered to ensure IP Service Activator is aware of the change.

During device discovery, if the discovered AS number differs from the default Domain AS number, IP Service Activator raises a warning against the device (Code 1627):

DeviceName (*DeviceManagementIPAddress*) ASN on device does not match Domain ASN

This warning is for information purposes only and does not inhibit service provisioning using the discovered AS number. If the device being discovered already has roles, due to re-discovery, this behavior is seen only for devices assigned the role of Gateway or Core. For devices assigned the role of Access, such as CE devices, AS number validation is not performed against the default domain AS. Instead, if the device is found to belong to a VPN site and if the site connectivity includes eBGP, the discovered AS number is matched against the peer BGP AS number specified on the VPN site.

If a managed device is re-discovered, and the newly discovered AS number does not match the previously discovered AS number, a warning is raised against the device (Code 1631): "".

The device *DeviceName* (*DeviceManagementIPAddress*) has a non-zero ASN in the OM that is different from the discovered ASN

The pre-existing AS number for the device with IP Service Activator is not changed. To actually change the AS number for the device, the device must first be unmanaged and then re-discovered.

IP Service Activator displays the discovered device AS number (ASN) in the Device dialog box on the **Device** property page.

Before Running Device Discovery

This section explains the steps that you need to take before starting the device discovery process:

- Ensure domain details are correctly set up.
- Ensure proxy agents are set up for automatic assignment.
- Specify the system's use of IP addresses for device management.
- Set up roles.

Setting Up the Domain Details

The way in which discovery works depends on whether public or private IP addresses are used within the domain. For the discovery process to work correctly, ensure that the domain type is set correctly on the **Domain** property page before running the discovery process. This can be **Public**, **Private** or **MPLS VPN**.

Note: You must create one domain per AS region.

See "Setting Up Domains" for information about setting up domain details.

Public Domains

If the domain is defined as Public, IP Service Activator can discover an entire network starting from a single device, because all IP addresses found are unique. The extent of the discovery can be controlled by specifying the number of hops to go from the original device. Alternatively, if no IP addresses or DNS names are known, a discovery can be started from the segment local to the policy server to find details of all connected devices, hosts and further segments.

A device can only be in one domain, but where a segment links two devices that appear in two different domains, the segment appears in both domains, as shown in Figure 5–1.

Figure 5–1 One Segment Appearing in Two Different Domains

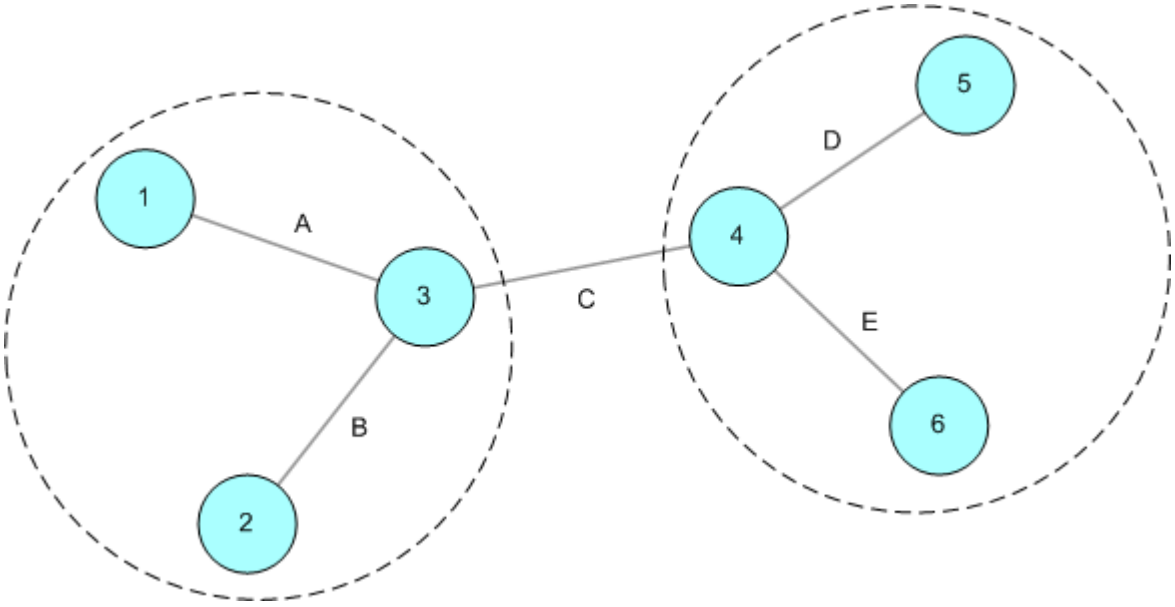

If devices 1, 2 and 3 are in one domain and devices 4, 5 and 6 in another, a representation of segment C appears in both domains.

In a public domain, the IP address of a device specified for the discovery process may be changed. For example, if a Cisco device has a loopback address defined, this is used to identify the device in preference to any IP address previously used. The default loopback address used is 0 but may have been configured for the domain.

See "Setting the Default Loopback ID Value for Discovery" for details.

Private Domains

If the domain is defined as Private, each device must be identifiable by a unique IP address, which may be a public address or a Network Address Translation (NAT) address. However, some or all of the interface IP addresses are likely to be non-unique. Therefore the discovery process requires a unique IP address or DNS name to be specified, and is not able to find further devices. The **Hops** field on the Discovery dialog box is therefore ignored.

Note: Connections between discovered segments are not automatically shown in the client, though these connections can be created manually by dragging one segment on to another.

MPLS VPN Domains

For MPLS VPNs, the core provider network is assumed to use public addresses, and IP Service Activator can discover an entire network starting from a single device. All CE routers are assumed to use private addresses and an IP address or DNS name must be specified in order to discover them.

If PE interfaces that are connected to CE devices have public addresses, the PE and CE devices will be automatically connected to segments during discovery, as shown in Figure 5–2. Depending on the role assignment rules applied, devices and interfaces will be assigned appropriate roles.

Figure 5–2 Public PE Interfaces connected to CE devices

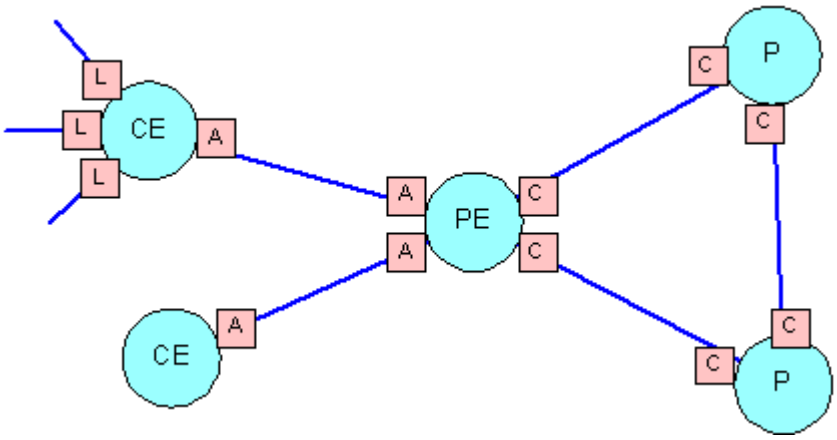

If PE interfaces that are connected to CE devices have private addresses, connectivity cannot be determined, as shown in Figure 5–3. Therefore CE devices and access interfaces on PE devices are not automatically connected to segments. In addition, because the devices are not directly connected, interfaces are not automatically assigned roles.

Figure 5–3 Private PE Interfaces Connected to CE Devices

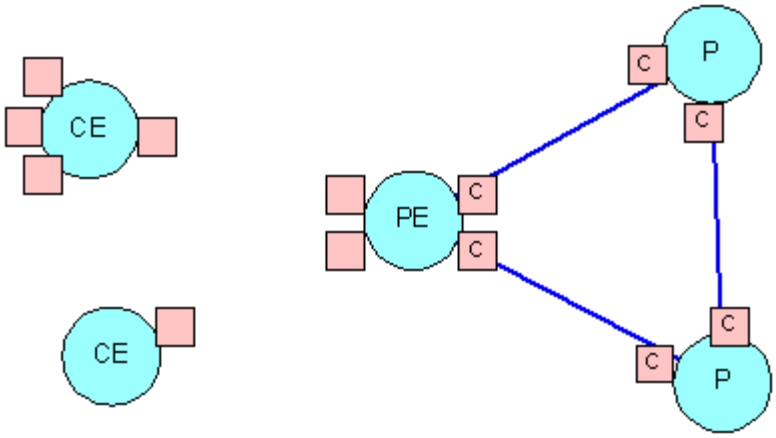

The connections between the PE and CE devices can be applied manually, by dragging one interface on to another on the topology map.

Note: If a role is assigned to the CE device, IP Service Activator does not show the CE device's connection to the PE device. If there is no role assigned to the CE device, however, IP Service Activator shows the connection.

Setting Up Proxy Agents for Automatic Assignment

All devices in the domain that are to be managed by IP Service Activator must be assigned to a proxy agent. It is the proxy agent that controls when and what type of configuration is to be applied to a specific interface.

Although it is possible to assign devices to proxy agents manually, it is generally performed automatically during device discovery.

See "Setting Up Proxy Agent Assignment" for information about setting up proxy agents for automatic assignment.

Defining How IP Addresses are Used

IP Service Activator uses a specific IP address when managing devices. Commonly this is the loopback address, but if this is not a unique address it is possible to use an alternative. On the Options dialog box you can set up a global option to define the standard way in which IP addresses are used. If necessary you can override this for individual devices (see "Setting Specific IP Addresses for Device Management").

Note: For complete dialog box and property page descriptions, refer to the Online Help.

To set global IP address usage:

1. From the **Tools** menu, select **Options**.

   The Options dialog box opens.

2. Select the **Discovery** property page.

3. Under **CE Device IP Address Selection**, select the global option used to set the IP address for CE devices. Choose from **Choose loopback**, **Choose first WAN address,** and **Do not change**.

   See "Setting the Default Loopback ID Value for Discovery" for more information.

4. Click **OK**.

Setting Up Roles

A role is a label that can be applied to one or more policy targets. Roles define:

- The policy or service configuration that is applied to that object
- For external collector systems, from which devices data is collected

See "Defining and Applying Roles" for information about defining roles.

Customizing Discovery Using AutoDiscovery.cfg

The **AutoDiscovery.cfg** file is used to specify how IP Service Activator interprets discovery information from devices in terms of translating the SNMP MIB-2 data into objects. It does this through a number of rules matching the formats shown below. Each rule defines a set of matching conditions. If a discovered item matches the conditions, it is reported into IP Service Activator as the type of object the rule specifies.

The **AutoDiscovery.cfg** file is located in `opt/OracleCommunications/IP Service Activator/Config`.

> **Note:** Be careful when editing the **AutoDiscovery.cfg** file. Do not change the settings unless you fully understand the impact on discovery. You should fully understand SNMP, ifTable, ifStack, and core MIB-2 concepts before creating changes to the **AutoDiscovery.cfg** file. It is important to construct these statements with extreme care, because conflicts are possible between statements and particularly, regex statements.

The Policy Server reads the **AutoDiscovery.cfg** file on startup. In order for changes to the **AutoDiscovery.cfg** file to take effect, the Policy Server must be restarted.

In general, the Boost regex (regular expression) is used.

The **AutoDiscovery.cfg** file is made up of the following types of entries. The basic mechanism is to evaluate each returned MIB-2 object from SNMP against the statements in the **AutoDiscovery.cfg** file. The discovery information is matched against the statement parameters to determine how IP Service Activator should handle the object.

The possible statement types in the **AutoDiscovery.cfg** are:

- Enterprise:`<enterpriseNumber>;<enterpriseName>;<deviceDriver>;<supported>;<accessType>;<configLevel>;`
- Subinterface:`<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;<Invert>;`
- Sublayer:`<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;<Invert>;`
- Vlan:`<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;<Invert>;`
- Vlanport:`<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;<invert>;`

- Main:<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;
- Ignore:<enterpriseNumber>;<highIfType>;<regexp>;
- Ifname:<enterpriseNumber1>;<enterpriseNumber2>;<ifType>
- Ifdesc:<enterpriseNumber1>;<enterpriseNumber2>;<ifType>
- Persistant:<enterpriseNumber1>;<enterpriseNumber2>;
- Volatile:<enterpriseNumber1>;<enterpriseNumber2>;
- Rename:<enterpriseNumber>;<ifType>;<pattern to match>;<pattern to substitute>;
- Icmp:[on | off]
- AtmVcInterfaceSource:<enterpriseNumber>;<aal5VccTable/atmVclTable>
- Host:<enterpriseNumber1>;<enterpriseNumber2>;
- Device:<enterpriseNumber1>;<enterpriseNumber2>;
- Controller:<enterpriseNumber>;<ifType>;<regexp>;

Most of the statements function in a similar way. Generally for statement X, if the MIB-2 information for an object matches the parameters in the statement, the object is treated as an object of type X in the IP Service Activator object model.

Enterprise

The Enterprise type defines a vendor to the discovery system. The Enterprise type identifies the device as belonging to a particular vendor and defines how IP Service Activator can communicate with the device, and what paradigm IP Service Activator uses when we configure the device.

The syntax is as follows:

Enterprise:<enterpriseNumber>;<enterpriseName>;<deviceDriver>;<supported>;<accessType>;<configLevel>;

where:

- <enterpriseNumber> - registered enterprise number to match on (e.g. Cisco = 9). This comes from the sysObjectId in the MIB2 report from the device for this vendor.
- <enterpriseName> - text string used to visually represent the vendor
- <deviceDriver> - the device driver / cartridge type that must be used to deal with devices from this vendor (if not supported, leave empty).
- <supported> - tells IP Service Activator whether any configuration other than discovery can be performed with this vendor's devices (yes/no).
- <accessType> - the default mechanism to get into the device for the purpose of command delivery (TACACS, NamedUser, anonymous, SNMPv1, SNMPv2c, none, PasswordOnly) - the details of the access mechanisms are documented in the Discovery section of the IP Service Activator online Help.
- <configLevel> - defines whether the device is to be configured as a whole or if it can be treated as a number of individual interfaces. Specify either Device or Interface.

For example:

```
Enterprise:9;Cisco;cisco;yes;TACACS;Interface;
```

This statement indicates that for devices with enterprise number 9, name the vendor as ‘Cisco', use the cisco device driver, access the devices using TACACS, and treat the device as though it has a number of interfaces below the device in hierarchy.

Subinterface

The syntax for Subinterface is as follow:

```
Subinterface:<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;<Invert>;
```

- `<enterprise_number>` - registered enterprise number to match on.
- `<highIfType>` - Interface IfType for higher layer in ifStack.
- `<lowerIfType>` - Interface IfType for lower layer in ifStack.
- `<regexp>` - regular expression to match the reported name on.
- `<Invert>` - invert flag. The optional Invert value indicates that IP Service Activator should invert whatever relationship was reported by SNMP. For example, if SNMP reports that 'B' is lower than ‘A’, we might actually want ‘A’ to be a Subinterface not ‘B’. The invert flag compensates for this situation.

The use of the Subinterface, Sublayer, and Main statements is to enforce the hierarchy of interfaces based on type (and regex matching on the name).

```
Subinterface:9;32;32;.*;
```

This statement indicates that an interface discovered that has an enterprise number of 9 and if ifStack reports a higher ifType of 32 and a lower ifType of 32, and the object is named anything, treat it as a subinterface.

Sublayer

The parameters for Sublayer are similar to Subinterface.

Sometimes devices deal with interfaces at a level of abstraction or granularity that is not reflected in the available discovery types in IP Service Activator. Objects can be reported as interfaces that are not actually configurable interfaces. For example, some artifacts of configuration are reported through SNMP as interfaces.

For example:

```
Sublayer:9;134;0;.*\.0;
```

The Sublayer statement is used to keep the reported object matching the statement in the IP Service Activator object model without storing it as an interface or subinterface.

Vlan and Vlanport

The syntax for Vlan and Vlanport are as follows:

```
Vlan:<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;<Invert>;
Vlanport:<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;<invert>;
```

The parameters for Vlan and Vlan port are similar to Subinterface.

VLAN and VLAN Port identify the discovered interfaces as VLAN or VLAN Port interfaces respectively as opposed to general use interfaces for devices for which these are not reported correctly through SNMP.

Main

The syntax for Main is as follows:

```
Main:<enterpriseNumber>;<highIfType>;<lowerIfType>;<regexp>;
```

The parameters for Main are similar to Subinterface.

For this an interface, if the enterprise number, interface pair, and optionally the regex match, treat it as a main interface in IP Service Activator.

For example:

```
Main:9;32;39;^Serial.*;
```

This example would map Cisco Frame Relay interfaces on a SONET controller as main interfaces in IP Service Activator if the description matches the regular expression ^Serial.*.

Ignore

The syntax for Ignore is as follows:

```
Ignore:<enterpriseNumber>;<highIfType>;<regexp>;
```

This means for the device matching the enterprise number, interfaces matching the high interface type number (and optionally matching the regex) are ignored. Discovery of matching interfaces will not be reported to IP Service Activator.

For example:

```
Ignore:9;94;.*-adsl;
```

This statement causes discovery to ignore ASDL interfaces on Cisco devices.

Ifname and Ifdesc

The syntax are as follow:

```
Ifname:<enterpriseNumber1>;<enterpriseNumber2>;<ifType>
Ifdesc:<enterpriseNumber1>;<enterpriseNumber2>;<ifType>
```

SNMP supports two different text representations of the interface - name and description. Some vendors put the canonical name of the interface in the 'name' field, while other put it in the 'desc' field. These statements allow you to account for these variations. By default, IP Service Activator uses the description to map to the description of the interface in the object model. However, by supplying an enterprise identifier and interface type in an Ifdesc statement, you specify that the discovered interface's name will be supplied from the description reported by SNMP.

For example:

```
ifname:9;5;0;
```

CatOS interfaces discovered store their name in 'name', not 'desc'.

Icmp

The syntax is as follows:

```
Icmp:[on | off]
```

Icmp controls whether IP Service Activator should ping a destination first before attempting SNMP discovery. This is used when a range of IP addresses (as specified by an address/mask subnet identification) are being discovered, or during segment-based discovery. All viable IP addresses are contacted and addresses which respond are discovered. This setting is 'on' by default.

Persistent

The syntax is as follows:

```
Persistent:<enterpriseNumber1>;<enterpriseNumber2>;
```

Volatile

The syntax is as follows:

```
Volatile:<enterpriseNumber1>;<enterpriseNumber2>;
```

For the given device (where its sysObjectId matches <enterpriseNumber1>,<enterpriseNumber2>), assume that the ifIndex values assigned to each interface are either preserved (i.e. Persistent) or recomputed on restarts or similar situations, and therefore can change (i.e. Volatile). This statement indicates to IP Service Activator whether the ifIndex value of an interface can be used to find matches and uniqueness. (When IP Service Activator re-discovers a device, it needs to map the interfaces coming out of SNMP discovery to the already-discovered interfaces in the object model).

Rename

The syntax is as follows:

```
Rename:<enterpriseNumber>;<ifType>;<string-to-match>;<string-to-replace>;
```

For this interface number, for this interface type, if you find an interface name that matches this pattern, use the substitute pattern to modify the name that is reported to IP Service Activator.

For example:

```
Rename:9;134;-atm subif;;
```

In this example, ATM subinterfaces are renamed so that the text "-atm subif" is removed from the name (i.e. is replaced with nothing).

AtmVcInterfaceSource

The syntax is as follows:

```
AtmVcInterfaceSource:<enterpriseNumber>;<aal5VccTable/atmVclTable>
```

There are two different SNMP MIBs that can supply virtual circuit information to IP Service Activator - the AAL5 table, and the ATM/VCL table. In some cases, SNMP reports data from both, even though only one is accurate. This statement lets you indicate, for the given enterprise number, which type of reporting of ATM VC Interfaces should be used by IP Service Activator.

For example:

```
AtmVcInterfaceSource:9;aal5VccTable;
```

Cisco devices will be have the aal5VccTable information reported to IP Service Activator for Cisco ATM VCs.

Host and Device

The syntax is as follows:

```
Host:<enterpriseNumber1>;<enterpriseNumber2>;
Device:<enterpriseNumber1>;<enterpriseNumber2>;
```

These statements let you indicate which enterprise number pairs indicate hosts, and which indicate devices, in terms of the IP Service Activator object model.

For example:

```
Device:3224;0;
```

Juniper Netscreen devices will be discovered as devices, not hosts.

Controller

The syntax is as follows:

```
Controller:<enterpriseNumber>;<ifType>;<regexp>;
```

Devices sometimes report hardware controllers as interfaces, even though they are not. This statement lets you indicate that reported items matching the values given in the statement are hardware controllers.

To discover controllers:

1. Take a backup of **Config/AutoDiscovery.cfg**.
2. Open **Config/AutoDiscovery.cfg** in a text editor.
3. Comment out the appropriate lines. For this, find blocks of lines delimited at the top by the line:

   ```
   //: Comment out the following lines to discover controllers
   ```

   The block is delimited at the bottom by the line:

   ```
   //: END Controller Discovery
   ```

   For each line in the middle of these two delimiters, add "//:" as a prefix to comment it out.

   Note: You can ignore lines that are already commented out.

4. Uncomment the appropriate lines. For this, find blocks of lines delimited at the top by the line:

   ```
   //: Uncomment the following lines to discover controllers
   ```

   and delimited at the bottom by the line:

   ```
   //: END Controller Discovery
   ```

   For each line between these two delimiters, remove the prefix "//:".

> **Note:** You can ignore lines that are commented out with the "//: " prefix. Note any additional space, these are really comments and should be kept commented out.

5. Save the **Config/AutoDiscovery.cfg**. Ensure that you have preserved user id, group id and permissions on the file. It is recommended you run a "diff" between the saved version and the modified version to confirm that all required changes have been made and no unwanted changes have been made.
6. Rediscover the devices.
7. All discovery operations initiated from now onwards will attempt to discover controllers as well.

Order of Evaluation

When comparing discovered information against the **AutoDiscovery.cfg** file, the Policy Server considers statements in this order.

1. Enterprise ID
2. Device or Host
3. AtmVcInterfaceSource
4. Interfaces: Main, Sub, Lay, VLAN, VLAN Port, or Controller
5. Ifname and Ifdef
6. Rename

> **Tip:** Use an SNMP MIB Browser to determine what information is coming from SNMP about your devices. Then you can determine which exceptions in the returned information you need to create statements for.

Running Device Discovery

After you set up the domain parameters and roles, you can set up the network topology. There are several alternative methods of doing this:

- You can initiate a discovery process from one or more device IP addresses or DNS names. This is the recommended method.
- If you do not know the names and addresses of any devices, you can initiate a discovery from the host system running the policy server.

> **Note:** Network discovery and related tasks, such as fetching capabilities, cannot be carried out part way through the current transaction. These tasks can only be performed immediately after a transaction has been saved or committed and before a new transaction is started.

To open the Topology Discovery dialog box, from the **Discovery** menu, select **Discover**.

The Topology Discovery dialog box opens.

Note: The **Discover** menu option is available only when network objects appear in the Details pane (either a map view or a details list).

The **Discover** menu option is not available if there are unsaved changes in the client.

Set up parameters on the three property pages of this dialog box:

- Discovery parameters specify the type of discovery
- SNMP parameters control the way SNMP works
- Security parameters specify access settings that apply to all discovered devices

Setting Up the Discovery Parameters

The settings you select on the Discovery property page depend on the type of discovery you want to perform.

Discovering One or More Specific Devices

If you know the IP address or DNS name of one or more of the devices in your network you can use the **Discover** option in the Topology Discovery dialog box to discover its details and the network nodes to which it is connected.

Figure 5–4 shows the Topology Discovery dialog box.

Figure 5–4 The Topology Discovery Dialog Box

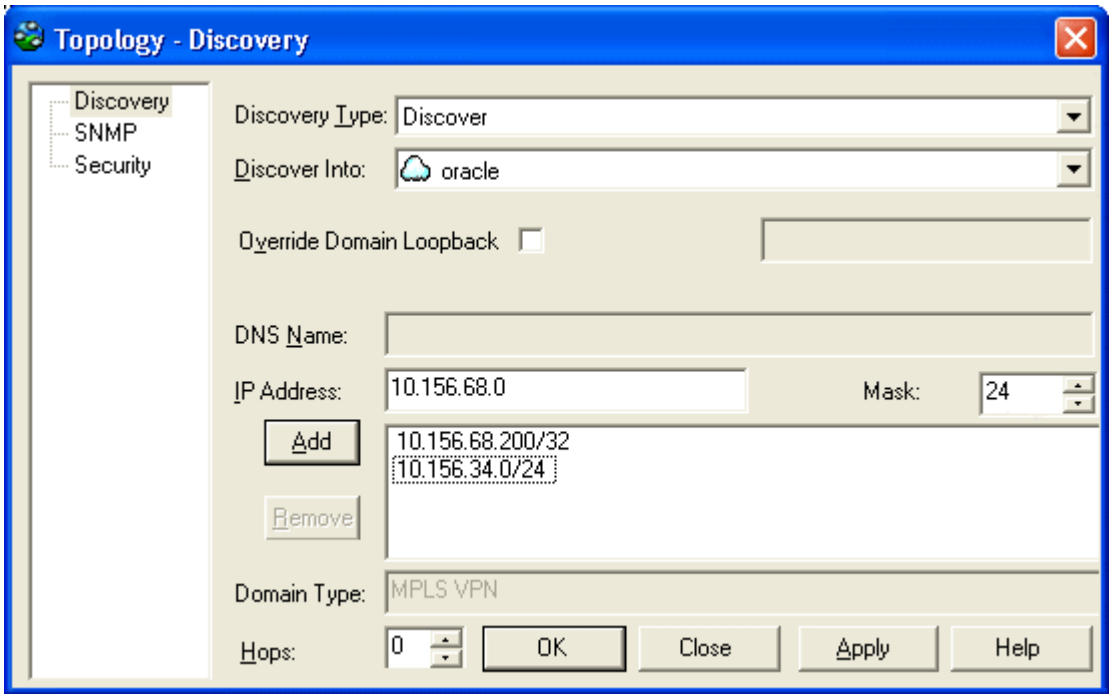

Note: A device can only be discovered into a single domain – a device cannot feature in two domains.

1. On the Topology Discovery dialog box, select **Discover** from the **Discovery Type** list.
2. In the **Discover Into** field, select the network to which you want discovered devices to be assigned.
3. (Optional) Override the default loopback ID value set for the Domain. See "Setting the Default Loopback ID Value for Discovery" for more information.
4. Enter the DNS name, or the IP address and mask for the device.
5. Click **Add** to add it to the list of devices to be discovered.

   The IP Address field accommodates both IPv4 and IPv6 addresses.

   The mask defaults to 32, unless an IP address in the format X.X.X.0 is entered in which it is defaults to 24. It can be set to any value from 16 through 32.

   The mask defaults to 32, unless an IP address in the format X.X.X.0 is entered in which it is defaults to 24. It can be set to any value from 24 through 32.

   You can discover a number of devices at once, either by listing them individually or by specifying an IP address and mask.
6. In a domain defined as Public or MPLS VPN, you can set the **Hops** field to a value between 1 and 10 to specify that the discovery process is to search for connected devices within a certain number of hops from the original device.

   The default hop number is zero, that is, only the specified devices are found. Increasing the hop count value broadens the discovery process and results in more of the network topology being discovered in a single discovery operation. Note that the more hops, the longer the process takes.

   Note: In a domain defined as Private the Hops field is ignored, as the discovery process requires a unique IP address or DNS name in order to find a device.

   It is not possible to discover Juniper M-series devices using the hop count method. These devices do not make their routing tables available via SNMP, the protocol used by IP Service Activator to interrogate devices and network segments.

   Figure 5–5 shows an example hop pattern from a starting device.

Figure 5–5 Example Hop Pattern from a Starting Device

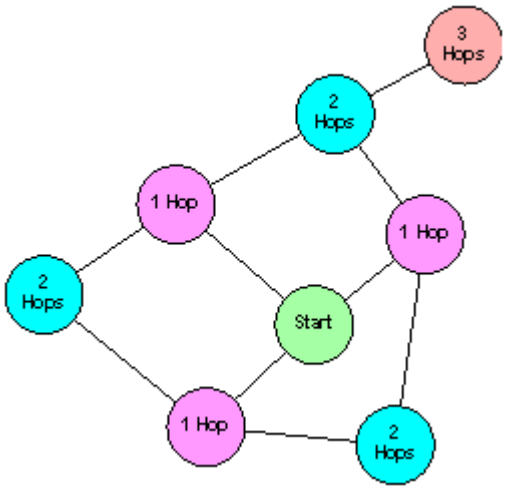

Discovering the Local Segment

If you do not know the IP addresses or DNS names of any devices in your network you can use the Local Segment option to start a discovery from the host system running the policy server.

> **Note:** This method cannot be used in a domain defined as Private. Note also that the policy server workstation must be running an SNMP agent.

To discover devices using the Local Segment option:

1. On the Topology Discovery dialog box, select **Local Segment** from the **Discovery Type** list.
2. In the **Discover Into** field, select the network to which you want discovered devices to be assigned.
3. Set the **Hops** field to a value between 1 and 10 to specify that the discovery process is to search for connected devices the specified number of hops from the original device.

   Because the segment on which this discovery process is performed includes the policy server, it is likely that the segment is a local subnet outside the enterprise or service provider WAN. Therefore, it is a good idea to set the hop value to at least three so that some nodes in the WAN are discovered.

Defining the SNMP Options for Discovery

The parameters controlling the way in which the discovery process operates are set up on the **SNMP** property page of the Topology Discovery dialog box.

On this property page, you can select **SNMP v1** and/or **SNMP v2c** to be used for device discovery. If you also want to use SNMP v3 for device discovery, you must select an **SNMP Profile** on this property page. More details follow:

- If you are using an SNMP Profile to set SNMP for discovery, select the check box and select a profile in the drop-down menu. An SNMP Profile allows you to use SNMP v3 for discovery, as well as SNMP v1 and SNMP v2C. Details for using SNMP profiles are provided in the IP Service Activator online Help.

IP Service Activator follows an SNMP "fallback" sequence until successful communication is established with a device. For example, if you select three SNMP versions for discovery, IP Service Activator starts the communication using SNMP v3, then falls back to SNMP v2c, and then to SNMP v1 if required, to establish the communication.

- The default number of **Retries** is 2 and the default **Timeout** period is 3 seconds. You may need to increase these values if your network is slow, or if you set a high hop count for the discovery process and select the **Next hop** route discovery method.
- By default both **SNMP v1** and **SNMP v2c** check boxes are selected. This means that the discovery process will try SNMP v2c first and, if this fails, will try SNMP v1. If you specify SNMP v2c only, there is a possibility that you will not discover some devices in your network. In particular, Cisco devices that are running an IOS earlier than version 12.0 do not support SNMP v2c. Selecting both versions of SNMP for the discovery process ensures that current and older router models can be discovered. If you are sure that only one version is in use you can deselect one version in order to speed up the discovery.
- The default **Read Community** setting is public; you need to change this if an alternative SNMP community has been set up on the devices.
- The **MaxRepetitions** field is used to set the value for max-repetitions in the get bulk operation when SNMP v2c is selected for discovery. It specifies the maximum number of rows that will fetch from a network resource in a single get-bulk request. The default value of this field is 100. Permitted range for this filed is 1 to 100. Decrease the value of Max-Repetitions if the routers don't respond to the default value of 100.

> **Note:** If the discovery of a Juniper device times out using SNMP v2 but can be successfully discovered using SNMP v1, then adjust this value.

- The **Discovery Method** setting is only applied if the **Hops** value on the **Discovery** property page is greater than zero. The default setting is **Next hop route**, which queries the device's routing tables – this is normally faster, but if you have Internet routers with very large routing tables, it can be slow. If you select **Segment** scan the discovery process scans each newly-discovered segment checking each address in turn. This process is comprehensive but can be slow.

When a device has been discovered, the specific SNMP settings used are saved as device-specific parameters. For example, if SNMP v1 was used, only the SNMP v1 check box is selected on the device-specific SNMP settings. The device-specific settings always override the global settings when devices are rediscovered.

> **Note:** If SNMP compression is configured on a device, discovery of that device will not work. Ensure SNMP compression is not configured on the device. Note that Juniper M-series devices have SNMP compression turned on by default.

Defining Default Security Options for Discovery

You are recommended to set up default security parameters before running the device discovery. This ensures:

- That security settings are automatically applied to all subsequently discovered devices.
- That QoS, security and VPN capabilities are retrieved from all devices (write access to devices is required in order to ascertain the capabilities).

If you do not set these values, you must set security settings for each device individually. In addition, IP Service Activator will be unable to obtain device capabilities.

The access parameters are set on the **Security** property page of the Topology Discovery dialog box.

Access Styles

The **Access Style** is the mode in which IP Service Activator accesses the discovered devices for configuration. Access Styles include:

- Named User
- Anonymous
- TACACS+
- SNMP v1
- SNMP v2c
- SSH with password authentication
- SSH with keyed authentication
- Password Only
- None

The selection determines which additional fields are displayed on the screen. Access methods are device-specific; at present, IP Service Activator configures devices as follows:

- Cisco devices are configured via the command-line interface using an anonymous user login with password authentication, via SSH with password authentication, or via a TACACS+ server if configured.
- Juniper M-series devices are configured via the command-line interface using Named User, via a TACACS+ server if configured, or via SSH with password authentication. (SSH with keyed authentication is not supported.)
- Juniper E-series devices are configured via the command-line interface using an anonymous login with password authentication.
- Alcatel devices are configured via the command-line interface using an anonymous login with password authentication, via a Named User login, or via SSH with password authentication. (SSH with keyed authentication is not supported.)

The access styles are as follows:

- For login as a Named User, specify the command-line interface login details (**Username** and device **Login** and **Enable** passwords).

> **Note:** This method is not supported at present.

- For login as an Anonymous user, you just need to set up the **Login** and **Enable** passwords.
- **TACACS+** indicates that a TACACS+ server is used to control access. The appropriate passwords must be set up at network level (see "Setting Specific Security Settings").
- For **SNMP v1** or **v2c** you need to specify the SNMP **Write Community** for write access to the device. This must match the community set up on the device(s).

> **Note:** This method is not supported at present.

> **Note:** SNMP v1 and v2c are not supported as access styles on any currently-supported devices.

- For **SSH with password authentication**, specify the **Username** and **Login Password**.
- For **SSH with keyed authentication**, specify the **Username** and select a **Key file**. Specify the private key file – the public key file must be present on the device.

> **Note:** For details on creating SSH keys for devices supporting this access style, refer to the appropriate Device Driver Support guide.

For the purposes of discovery, you can leave the access style as None. In this case, IP Service Activator applies a default access style according to the device type. The access styles are defined in the **AutoDiscovery.cfg** configuration file, which is installed in the **Config** subdirectory of the host system running the policy server. However, login and password information cannot be set and therefore must be set manually for each device.

> **Note:** Note that the specified access settings apply as defaults. You can amend them for specific discovered devices, in which case the device-specific settings will override the global settings when devices are rediscovered.

Cartridge Support for Access Styles

Table 5–1 displays the Access Style supported for each cartridge, using the Network Processor. **Y** indicates the style is supported, whereas **N** indicates it is not supported.

Table 5–1 Cartridge Access Styles

Access Style	Huawei	Cisco	Juniper	Brocade
Named User	N	N	Y	N
Anonymous	Y	Y	N	N
TACACS+	Y	Y	Y	Y
SNMP v1	N	N	N	N
SNMPP v2c	N	N	N	N

Table 5–1 (Cont.) Cartridge Access Styles

Access Style	Huawei	Cisco	Juniper	Brocade
SSH with password authentication	Y	Y	Y	Y
SSH with keyed authentication	N	N	N	N
Password Only	N	Y	N	N
None	N	N	N	N

The Discovery Process

As soon as the changes are committed to the database, the network discovery starts. IP Service Activator initially creates a list of nodes to be discovered. When discovering a segment, each additional node found is added to the end of the list. If a timeout or failure occurs, the process moves on to the next node on the list.

Discovered devices are linked to the network object that was selected when the discovery process was started.

After discovery, devices are allocated to a proxy agent (see "Setting Up Proxy Agent Assignment"). Where possible, device and interface capabilities are ascertained.

You must also assign roles to devices and their interfaces.

Monitoring the Discovery Process

The length of time that the discovery process takes varies depending on the number of devices and segments found, the type of discovery and the speed of the network. If you have set a high hop count, a large number of devices may be found.

The status bar indicates the number of discovered devices.

The number indicates the number of nodes currently on the discovery list. This can initially increase as new devices and segments are found and interrogated, and will return to Idle when the discovery is complete.

When discovery is complete, IP Service Activator interrogates each device for its interface capabilities. The status bar indicates that this phase is in progress.

If a device's capabilities could not be obtained, or IP Service Activator obtained capabilities for one but not all of its interfaces, an error is generated.

The topology section of the object model is updated with details of the discovered devices, segments and hosts.

Automatic Mapping

By default, you create a network map by arranging devices manually. However, if you have selected automatic layout, the discovered network objects are added to the network map automatically while the discovery process runs.

In automatic mapping, you create a layout filter that defines which network objects are displayed on the map. Up to 200 network objects can be mapped automatically with minimal delay in display time. For greater numbers of network objects, there may be some delay in mapping and displaying the objects.

See "How Objects are Represented" for information on automatic and manual mapping options.

Stopping the Discovery Process

If necessary, you can stop the discovery process.

To stop discovery of the current object only:

- Select **Stop** from the **Discovery** menu

The discovery process continues with the next node on the list, if there is one.

To terminate the discovery process completely:

- Select **Stop all** from the **Discovery** menu or select **Stop all** from the **Discovery Type** drop-down list on the Topology Discovery dialog.

All devices and segments discovered prior to this command are retained in the object model, but discovery queries in progress at the time are stopped and the rest of the items on the discovery node list ignored.

The ERX Discovery Module

Juniper ERX devices are discovered using the ERX Discovery Module, which is installed along with IP Service Activator. The ERX Discovery module can be accessed through the client.

> **Note:** In the Configuration GUI, ensure that the **Juniper JUNOSe Discovery Module** check box under Component Manager Entries is selected. See *IP Service Activator System Administrator's Guide* for more information on the Configuration GUI.

To discover Juniper ERX devices:

1. Right-click the domain into which you want to discover the devices.
2. Select **Modules** from the context list, and then select **ERX Discovery**.
3. Select **Discover New ERX**. This displays the ERX Discovery dialog box.
4. In the ERX Discovery dialog box, select **Discovery**. This displays the Discovery property page.
5. In the Discovery property page:
   - Select the **Discovery Type** and the domain into which you want to discover the Juniper ERX device
   - Enter the IP address of the ERX device you want to discover.
   - Select the device from the Device Type list
6. In the SNMP property page:
   - Enter the read community value
   - Enter the values for the number of retries and timeouts
7. In the Security property page, enter access-related information in the Security property page
8. Click **OK** to start the discovery process.

After Discovery is Complete

This section explains the steps you need to take after discovery is complete to ensure your network is set up satisfactorily. You need to check the following:

- Ensure devices are assigned policy roles, and apply roles manually if necessary.
- Ensure devices are assigned to proxy agents.
- Set any specific IP addresses used to manage devices.
- Define how the system is to deal with manually-applied configuration.
- Set any specific security settings required for IP Service Activator to configure devices.
- Create your MIPSA registry in the cartridge tool (Cisco).
- Set devices to Managed to ensure that they will be configured by IP Service Activator.

You should also check any errors listed in the current faults pane and correct any problems. Occasionally, a new device type may be reported.

Ensuring Devices Are Assigned Roles

See "Assigning Roles" for information on applying roles manually.

Ensuring Devices Are Assigned to Proxy Agents

All devices that are to be managed by IP Service Activator must be assigned to a proxy agent. This is normally done automatically during device discovery, but if devices are not assigned to the correct proxy agents, you need to assign them manually.

To check which devices are assigned to a proxy agent:

1. On the **System** tab, open the **System Hosts** folder.
2. Open the component manager and the relevant proxy agent. IP Service Activator lists the devices that are assigned to the proxy agent in the Hierarchy pane.

To check whether a device is assigned to a proxy agent:

1. Display the ancestry pane by selecting **Ancestry Tree** from the **View** menu.
2. On the **Topology** tab, select the device.

If the device is assigned to a proxy agent, IP Service Activator lists the proxy agent in the ancestry pane.

To assign a device to a proxy agent manually:

1. In the Details pane, double-click the relevant network cloud (display devices in Map view or Details view)

   The devices are displayed.
2. Select the **System** tab and display the proxy agent in the System Hosts folder.
3. Drag the device from the Details pane and drop it on to the proxy agent that will control it.

   Alternatively, select the device and select **Copy** from the **Edit** menu (or click the **Copy** button). Then select the proxy agent on the **System** tab and select **Paste Link** from the **Edit** menu.

To unassign a device from a proxy agent:

1. Select the **System** tab and display the proxy agent in the **System Hosts** folder.
2. Select the device under its parent proxy agent and click the **Unlink** button on the toolbar.

Setting Specific IP Addresses for Device Management

IP Service Activator requires a unique IP address when communicating with and managing devices. This is frequently a loopback address, but an alternative can be specified.

You can set up global standards for IP address usage in the Options dialog box (see "Defining How IP Addresses are Used"). However, if necessary you can override these defaults with device-specific settings.

> **Note:** For Juniper E-series devices, the address used to manage the device must not be changed from that used for discovery. This is because the IP address used must be the one valid for the default virtual router. If an alternative address is used, the device is discovered but the driver cannot communicate with it.

To set a specific IP address for device management

1. Right-click the device, and select **Properties**.

   The Device dialog box opens.
2. Select the **Management** property page.
3. Select the IP address to be used for managing the device by choosing from one of the known addresses listed in the **Management IP Address** list, or by entering the management IP address in the edit box.
4. Click **OK**.

Setting Manual Configuration Detection

You can specify the action that IP Service Activator takes if it discovers that a device has been configured manually. A domain-level setting specifies the default behavior (see "Setting Up Domains"), but you can override this for particular devices.

> **Note:** This section does not apply to the Network Processor.

To set up manual configuration behavior for a device:

1. Right-click the device, and select **Properties**.
2. Select the required behavior from the **Manual config** list:
   - **Inherit from domain**: the setting for the device is inherited from the domain-level setting. For more information, see "Setting Up Domains".
   - **Delete**: if manual configuration is detected, it is deleted and the device reset to the original configuration.
   - **Warn and delete**: if manual configuration is detected, a warning message is output and the device is reset to the original IP Service Activator configuration.

- **Fail and don't delete**: if manual configuration is detected, a critical fault is raised. The device status is set to Intervention Required and no configuration is applied.

> **Note:** For VPN-related configuration, IP Service Activator always preserves manually pre-configured VRF tables, even if a delete option is selected in the **Manual config** field. However, BGP configuration is deleted unless the **Fail and don't delete** option is selected.

Setting Specific Security Settings

Before you can configure the network, you need to set up security and authentication settings to allow IP Service Activator to access the managed devices.

If you set up security values before you run a device discovery, (see "Defining Default Security Options for Discovery") these settings apply to all discovered devices within the domain. However, you can overwrite these with device-specific settings. If you have not set default security settings, you need to set suitable values for all devices to be managed.

The security settings you need to define depend partly on:

- The requirements of the device driver.
- The set-up of your network – for example, whether you are logging on as a named user or authenticating via a TACACS+ security server.

If you are using the same authentication details for all devices, you can set the security parameters at network level. If you have created subsidiary networks within the domain, you can specify that these networks inherit access parameters from the parent network.

To set security parameters for a network:

1. On the **Topology** tab, right-click the network object, and select **Properties**.
2. Select the **Security** property page.
3. (Optional) Click the **Inherit from parent network** check box.

   Do this if it is a sub-network. The security settings are then inherited from the network to which the sub-network belongs.
4. Enter the details required.
5. Specify the private key file for SSH with keyed authentication.

   A public key file must be present on each device in the network

> **Note:** Note that if the **Inherit from parent network** check box is selected, these field values are inherited from the parent network object and are therefore read-only.

6. Click **OK** to close the dialog box.

To check or set the security parameters for a device:

1. From the Topology tab, right-click the device, and select **Properties**.
2. Select the **Security** property page.

3. To set individual security parameters for this device, deselect the **Inherit network login** check box.
4. Select an Access Style and set parameters.
5. Click **OK**.

Managing Devices

Managing and unmanaging a device controls whether or not IP Service Activator configures the device.

A device's status is automatically set to Unmanaged when it is first discovered. Before a device can be configured, its status must be set to Managed.

> **Note:** A device must also be assigned to a proxy agent before it can be configured by IP Service Activator, even if its status is set to Managed. See "Assigning Devices to Proxy Agents" for more information.

You can stop a device from being managed by setting its status to Unmanaged. An Unmanaged device's configuration is not updated when you commit a transaction. You should unmanage devices that are no longer within the policy domain. You may also need to unmanage a device temporarily if attempts to configure it with IP Service Activator result in a fault that cannot be resolved.

You can manage or unmanage devices on a domain-wide basis or device-by-device. When you manage or unmanage a device, its color changes to reflect its new status. A managed device is represented by a green icon, an unmanaged device is represented by a blue icon.

To manage or unmanage all devices within a domain:

- Right-click the network object that represents the entire domain, and select **Manage All Devices** or **Unmanage All Devices**.

To manage or unmanage an individual device:

1. Right-click a device, and select **Properties**.
2. On the Management property page, click the **Manage** or **Unmanage** button.

Retaining or Removing IP Service Activator Configuration

By default, when you unmanage a device any configuration that IP Service Activator writes to the device remains configured. However, you can specify that the configuration is removed when a device is unmanaged. You can do this on a system-wide or device-by-device basis.

To specify the unmanage action globally:

1. From the **Tools** menu, select **Options**.

   The Options dialog box opens.
2. Under **Unmanaged Action**, select **Leave**.

   The IP Service Activator configuration is maintained on devices when they are unmanaged.

To specify the unmanage action for a device:

1. Right-click a device, and select **Properties**.

   The Device dialog box opens.

2. Select the **Management** property page and select one of the following:

   - **Inherit System Action**: applies the global unmanage action to this device (this is the default setting).
   - **Leave**: maintains the IP Service Activator configuration on the device when it is unmanaged

     Selecting **Leave** overrides any global setting.

Discovering a New Device Type

A device type defines the device driver to use to control the device. Every router in the network is classified according to type. The association between a router type and the device driver that will be used to control it is defined in the **DeviceTypes.cfg** configuration file. The information from this file populates the **Device Types** folder on the **Topology** tab.

Figure 5–6 shows the device types list in the Hierarchy pane.

Figure 5–6 The Device Types List

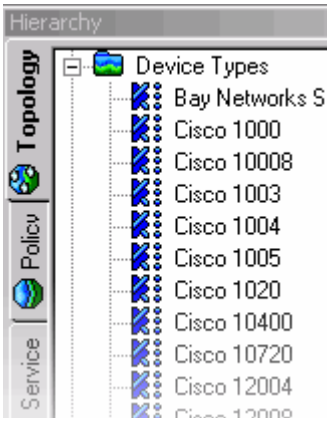

Some of the routers in the network may fall outside the classification types recognized by IP Service Activator. Where this is the case, IP Service Activator classifies the router according to its SNMP SysObjectId parameter during the discovery process.

For details of the devices currently supported by IP Service Activator, see the relevant device driver guide.

When the discovery process runs, IP Service Activator checks the device type of each discovered device against the **DeviceTypes.cfg** file. If it discovers a new device type, IP Service Activator automatically adds its details to the database and displays notice 1609 - New device type has been created in the current faults pane. This indicates that you need to check and amend the type's details and set up the appropriate device driver in the Device Type dialog box

Note: You may discover some devices that cannot be managed by IP Service Activator. The discovery process discovers most SNMP-enabled devices and hosts that support MIB-II. If you discover a device type that is not supported by any device driver, you can configure the new device type's details but attempts to configure the device are likely to result in error messages.

To set up a new device type:

1. Double-click the relevant message 1609 in the current faults pane.

   The Device Type dialog box opens.

2. In the **Device driver** field, enter the name of the device driver that you want to use to manage this device. This must match exactly the name specified in the appropriate Driver properties dialog box. You can access a device driver's properties dialog box via the **System** tab and the **System Hosts** folder.

3. (Optional) Change the **Vendor** field and **Product** field to provide more meaningful information.

   These fields are for information only.

4. The **Model**, **s/w Vn** field identifies the device type, and appears on the Device property page, so it can be useful to provide more meaningful information. To edit this field, select the **Edit fields** check box.

   Caution: Do not edit the **SysObject ID** field and **Configure level** field.

5. Click **OK**.

   The data is saved when the transaction is committed.

   Caution: If you set up a new device type in this way, it is at your own risk. If you require this device to be certified as tested by Oracle Communications, you must contact Oracle Communications Technical Support.

Reducing the Number of Listed Device Types

The **Device Types** folder lists a large number of device types. If you delete a device type from the folder, it is recreated the next time the policy server is started because the **DeviceTypes.cfg** file is reloaded at this point. If you want to reduce the number of listed device types, you can edit the **DeviceTypes.cfg** file using a text editor.

Caution: We strongly recommend that you create a backup copy of the file before editing.

Managing Configuration Thresholding

Use configuration thresholding to restrict the number of commands provisioned on devices in a single transaction.

About Configuration Thresholding

Configuration Thresholding provides a safety mechanism that blocks any device configuration action by IP Service Activator that exceeds certain user-specified parameters. The threshold is configured by means of two values: a regular expression (regex) against which to match commands, and the threshold value itself.

The regular expression uses a syntax based on the Boost regex library. See the Boost Website at:

http://www.boost.org/doc/libs

Setting the Configuration Threshold

This feature is turned off by default. All Network and Device settings for the threshold are set to **No Limit**. The feature starts working when you change this setting for a Network or Device.

The regular expression against which to match commands is not accessible through the client. The threshold value can be set at the **Network** or the **Device** level.

Configuration Threshold settings applied at the network level (on the Network Properties dialog box, **Device Management** property page) are automatically inherited by the other networks and devices inside that network.

These settings can be overridden for individual devices (or networks under a parent network for that matter.) Use the Device Properties dialog box, **Management** property page to set values for a device. In most cases, you will configure devices to inherit values from their parent network. Only the most sensitive devices (such as PE devices) require specific settings.

If you make changes, you must wait for propagation to occur before they are enforced.

See "Setting up Configuration Thresholding" for step-by-step instructions.

What Happens When the Threshold is Exceeded

If a transaction causes more configuration commands than specified by the threshold value, which matches the regular expression, the Telnet session to that device is ended and everything done in that Telnet session is undone. No commands are sent to the device, and the device is placed into Intervention Required mode.

As well, a Critical fault is raised against the router (visible in both the client and the OIM). This fault provides the threshold value as well as the current match count in the fault details. The aborted commands are logged to the Audit Trail file for that day with the prefix "max-transaction-exceeded".

The threshold parameters are applied only to commands generated by the IP Service Activator device driver. Specifically, it is not applied to commands generated by CDK scripts and modules.

Recovering from an Exceeded Threshold

When a device enters the intervention required state, no further configuration is sent to the device. To recover, you must examine the aborted commands in the audit trail. If these commands are legitimate (that is, it was your intent to issue such commands), you must increase the threshold value for that router, delete the fault, and commit.

If the commands are incorrect due to operator error, you must undo the actions that led to the behavior (for example, taking a role off a device), delete the fault, and commit.

If the commands are incorrect due to an IP Service Activator internal problem, or if you are not able to make a determination, contact Global Customer Care.

It is highly recommended that you undertake device recovery during a maintenance period.

Configuration Thresholding Notes

The command pattern to match against (that is, the regex) can be viewed, but not modified, in the client. It can be modified only through the OIM.

Some features with deficient implementations unnecessarily remove and re-install configuration on devices. Such removal commands may be counted towards the maximum transaction size, if they match the pattern.

If the threshold value is modified and the role removed from a device in the same transaction, the new threshold does not take effect until a new role is added to the device.

If the threshold value is modified and the device unmanaged (with the unmanage action set to remove) in the same transaction, the new threshold does not take effect until the device is re-managed.

Setting up Configuration Thresholding

This topic contains a sample procedure for setting up Configuration Thresholding.

> **Tip:** It is recommended that you group devices that are at the same functional category. For example, PE routers that behave the same way can be grouped within sub-networks.

To set configuration thresholding values:

1. Right-click on a network, select **Properties**, and set the maximum number of removal commands per device session on the Network Properties dialog box, **Device Management** property page.

   > **Note:** By default, devices and sub-networks inherit the configuration thresholding settings from their parent network.

2. Enter values for your network configuration, including **Inherit from Parent Network**, **No Limit**, **Limit To**, and **Match Expression**.
3. If there are other sub-networks, repeat the previous step to set the threshold limit for each of them.
4. Customize the settings for a devices by right-clicking the device, selecting **Properties**, and accessing the device properties - Management property page. Uncheck the **Inherit from network** check box, and provide device-level settings.

Setting the Configuration Thresholding Regex Match Expression

> **Note:** The commands included in the threshold count can be set only using Integration Manager or using OIM python scripts.

To use the Integration Manager, launch its CLI and log in.

If you know the name of the network to set up, skip this step. To find all the networks and their IDs, use the command:

```
find / network:"*"
```

To set the regular expression to match configuration statements when counting for Configuration Thresholding for a particular network:

```
find / network:"<network-name>"
getattributes <network id>
modify <network id> MatchesPatternTransactionSize="<regex>"
commit
```

A sample regex expression is: `no\\s(?!alias|auto-summary|synchronization)`

This would count any configuration statements starting with the word "no" except for those that start with "no alias", "no auto-summary" or "no synchronization".

Note: The threshold value can be set using `MaxTransactionSize` or set using the client.

Creating Virtual Devices and Interfaces

As well as discovering existing devices, you can manually create device and interface objects to represent those that IP Service Activator has not discovered. This allows devices not managed by IP Service Activator, such as those in the core network, to be modeled in the system in the form of virtual devices.

Note: Virtual devices and their interfaces cannot be discovered or managed.

To create a virtual device:

1. Right-click a network object, and select **Add Device**.

   The Device dialog box opens.

   Note that the Management, Security, Script Context, Capabilities, and Discovery property pages are not displayed.

2. Enter identifying details for the device.

   You can set the **Name**, **IP Address**, **Device Type**, and **Description** fields; other information is read-only.

3. Click **OK**.

Devices created in this way always have a status of Virtual, and the device icon appears in dark gray.

To create an interface on a virtual device:

1. Right-click the virtual device object, and select **Add Interface**.

   The Interface dialog box opens.

   Note that the Capabilities, Script Context, and Details property pages are not displayed.

2. Enter identifying details for the interface.

You can set the **Number**, **IP Address**, **Description** and **Media Type** fields; other information is read-only.

3. Click **OK**.

Interfaces created in this way always have a status of Unknown, and the interface icon appears in dark gray.

> **Note:** You can set roles for virtual devices and their interfaces, but this is only for identification purposes because they are not managed or configured by IP Service Activator.

Maintaining the Network Topology

You need to ensure that the network you are managing is accurately represented within IP Service Activator. For example, if new routers are added or existing ones reconfigured, you must ensure that IP Service Activator is updated accordingly. You can rediscover the entire domain, update individual devices and if required, set up the discovery process to run automatically on a regular basis.

Refreshing the Entire Domain

You can rerun the discovery process at any time if changes are made to the network configuration, or to ensure that the information you have is accurate.

> **Note:** The rediscovery of an entire domain may take some time. You can monitor its progress from the status bar at the bottom of the screen.

To update the topology of a domain:

1. Within a domain, select **Discover** from the **Discovery** menu.

   The Topology dialog box opens.

2. On the Discovery property page, select **Domain Refresh** from the **Discovery Type** list.

3. (Optional) Set the **Hops** field.

> **Note:** The Hops setting only applies to new devices and segments found.

4. Click **OK**.

> **Note:** You cannot change the SNMP or Security settings on a domain refresh. The settings that were saved in the database from the previous discovery for each device or host are used.

Rediscovering Individual Devices

You can re-run the discovery process at any time for a selected device, segment, or host.

To re-run device discovery for an existing device, segment, or host:

1. Select the device, segment, or host (either on the map or anywhere within a domain management window).
2. From the **Discovery** menu, select **Refresh**, or right-click the object and select **Discover**.

   IP Service Activator uses SNMP to discover details of the device, such as its interfaces, or any further devices or segments connected to it.

   New nodes that are discovered appear in the palette.
3. (Optional) Drag new nodes on to the map and run the discovery process on each new node in turn.

Deleting Missing Interfaces

The IP Service Activator discovery process uses SNMP to identify devices and interfaces and stores a representation of the discovered network topology in its internal object model. Interfaces are discovered by their ifIndex values, but IP Service Activator also records the value of each interface's ifName and/or ifDescription variable.

On a rediscovery of the network, previously discovered interfaces are primarily matched against interfaces encountered in the current discovery using the interface name (based on the ifName/ifDescription). However, if no match can be made, IP Service Activator compares the interfaces' ifIndex values and, where values coincide, matches the interfaces.

If IP Service Activator cannot match details held for an interface in its object model against any ifName/ifDescription or ifIndex values encountered during a new discovery, it regards the interface as 'missing'. A missing interface is assigned a status of Not Found and its ifIndex value is set to zero in IP Service Activator's object model.

Note that an interface is only classified as missing if it has been assigned a role or is linked to a site in a VPN or a layer 2 site in a TLS. If neither of these conditions apply to a missing interface, the interface is automatically deleted.

An interface whose status is **Not Found** is represented in the following way in the client.

Possible reasons for an interface to be assigned the Not Found status are:

- The interface has been deleted from the device, either deliberately or accidentally.
- The interface has been relabelled on the device – that is, both the interface's ifName/ifDescription and ifIndex details have changed. This sometimes occurs, for example, when a device's operating system is upgraded.

> **Caution:** You must investigate the cause of an interface being reported as Not Found before you consider deleting the interface.

If the Interface is Deleted

The interface might be deleted deliberately or by accident:

- If the interface is deleted deliberately, delete the Not Found interface in IP Service Activator.
- If the interface is deleted accidentally, re-create the interface on the device (you can Telnet to the device from the client) before rediscovering the device.

When a device that has 'not found' interfaces is unmanaged, the client displays a popup warning about possible loss of configuration.

Neither the client nor the OSS Integration Manager interface allows you to manage a device with 'not found' interfaces. These have to either be discovered to be found or must be deleted individually.

Interfaces which are 'not found' can be restored to 'found' state by using the preserve missing interfaces command:

To preserve all missing interfaces on a device, right-click the device and select **Preserve Missing Interface(s)**.

IP Service Activator restores the interfaces to its representation of the network.

To delete missing sub-interfaces on an interface, right-click the device and select **Preserve Missing Interface(s)**.

IP Service Activator restores the missing sub-interfaces to its representation of the network.

To preserve a specific missing interface or sub-interface, right-click the interface or sub-interface and select **Delete**.

Interfaces can also be preserved from the OSS Integration Manager interface. Issue the command:

`preservemissinginterfaces <object_id>`
where `<object_id>` is the identifier for the device, interface, or sub-interface.

Note: If an Interface or Sub-Interface is deleted when it is in Offline maintenance mode (either because the parent device, parent interface, sub-interface, or network processor is in Offline maintenance mode) the configuration on that device will not be removed from the device, but IP Service Activator will lose all knowledge of that configuration. In addition, if the deleted interface or sub-interface was created by IP Service Activator, it also will not be removed from the device, but IP Service Activator will lose ownership of that interface/sub-interface. The interface/sub-interface will reappear in the client during the subsequent re-discovery of that device.

Setting Up Discovery to Run Automatically

Because a full rediscovery of a large domain can take some time, you can set up IP Service Activator so that the device discovery process is run automatically every day, for example as an overnight process.

To set up discovery to run automatically:

1. Select the **System** tab on a global or domain level window.
2. Right-click the policy server object (master_server) and select **Properties**.
3. Select the **Discovery** property page on the Policy Server dialog box.
4. Click the **In background** check box and enter the time at which you want the discovery process to start each day
5. Click **OK**.

The discovery process is controlled by the SNMP parameters on the Discovery property page for each node. You should ensure these are set up correctly.

After the network discovery process has run each day, you should check the network topology and make any necessary corrections. For example:

- If a device is not found, it is not deleted, but its status is changed to Not Found. You will need to check it.
- If you are using manual mapping, any newly-discovered devices and segments are listed in the palette. You can drag them on to the appropriate map and ensure they are set up correctly.
- Messages relating to new or incorrectly-configured devices may be reported in the current faults pane. You should fix any problems before proceeding.

6

Representing and Mapping Objects

The Oracle Communications IP Service Activator client provides a graphical and hierarchical representation of the network objects included within your policy domain. You can create multiple maps that allow you to customize how information is displayed.

This chapter:

- Describes how objects are represented in the client, including VLANs and virtual routers
- Explains how to create a topology map manually or automatically
- Provides guidelines for working with maps, including recalculating map layout and configuring the palette
- Explains how to create subsidiary maps and alternative map views
- Describes how to display background images and change a map's scale

How Objects are Represented

After a network is discovered, details of devices, interfaces, segments and so on appear in the **Topology** tab of the Hierarchy pane and in the Details pane. See "IP Service Activator Windows" for details about how this information is displayed.

Alternative Device Views

You can choose two alternative views of devices and interfaces on the client using toolbar buttons:

- **Status Context View**: all devices and interfaces are shown color-coded according to their status.
- **Policy Context View**: all devices are shown in shades of gray according to their policy role.

Status Context View

If you click the **Status Context** icon, different colored icons are used to represent the different status values that a device or interface may have. Figure 6–1 indicates which status each icon represents and defines each status.

Figure 6–1 Status Context Icons

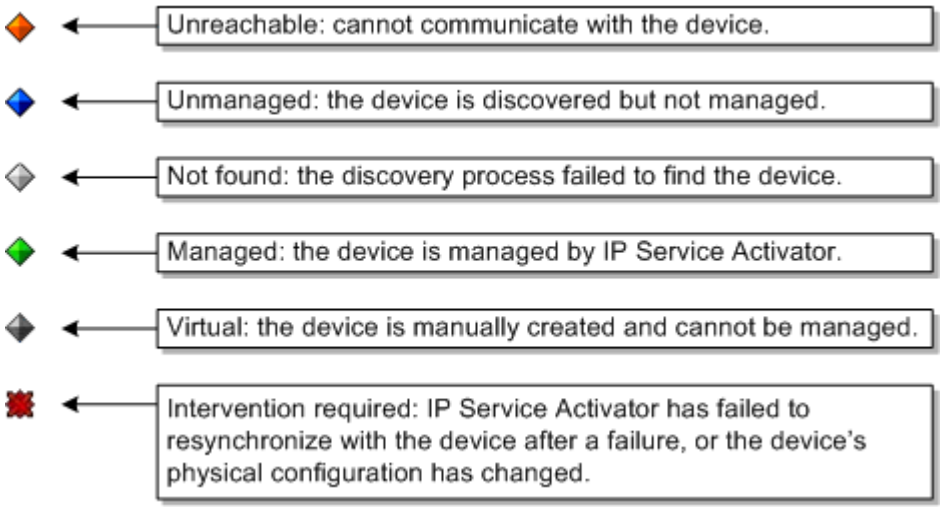

Policy Context View

If you click the **Policy Context** icon, different shadings are used to represent different system-defined roles of devices and interfaces. Figure 6–2 indicates which role each shading represents.

Figure 6–2 Policy Context Icons

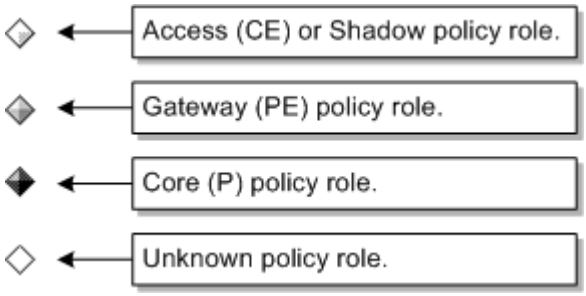

In the policy context, different colors are used to represent the different policy roles of links (that is, the connected interfaces):

- Access links are shown in Blue
- Core links are shown in dark blue
- Local links are shown in light blue
- Disabled links are shown in dark gray
- Links classified as None are shown in light gray
- Links in conflict, where the roles assigned to either end of the connection do not match, are shown in purple

Object Labels

By default, each network component displayed on the map is labeled with the object's name, description, or net address. The label used depends on the object type. If you have assigned system-defined roles to devices, the role for each device is also displayed.

To switch object labels on or off:

- From the **Map** menu, select **View Labels**.

  Object labels are displayed on the map when a check is displayed next to the **View Labels** menu option.

Network Segments

You can choose to hide the connectivity of a network segment if it is not shown on the map. This is controlled by the **Hide connectivity when not on map** check box on the Segment property page. By default this check box is deselected, and objects connected with a segment are always shown as connected, whether or not the segment appears on the map. If this check box is selected, no connectivity is shown between objects that are indirectly connected by the segment.

Figure 6–3 shows the connectivity of a network with the **Hide connectivity when not on map** check box disabled.

Figure 6–3 Hide Connectivity When Not on Map Check Box Disabled

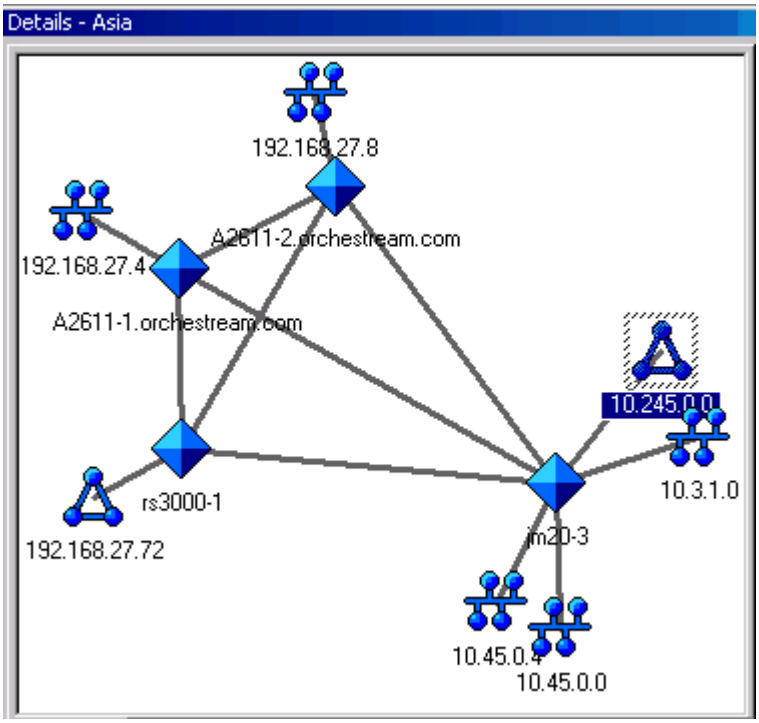

Figure 6–4 shows the connectivity of a network with the **Hide connectivity when not on map** check box enabled.

Figure 6–4 Hide Connectivity When Not on Map Check Box Enabled

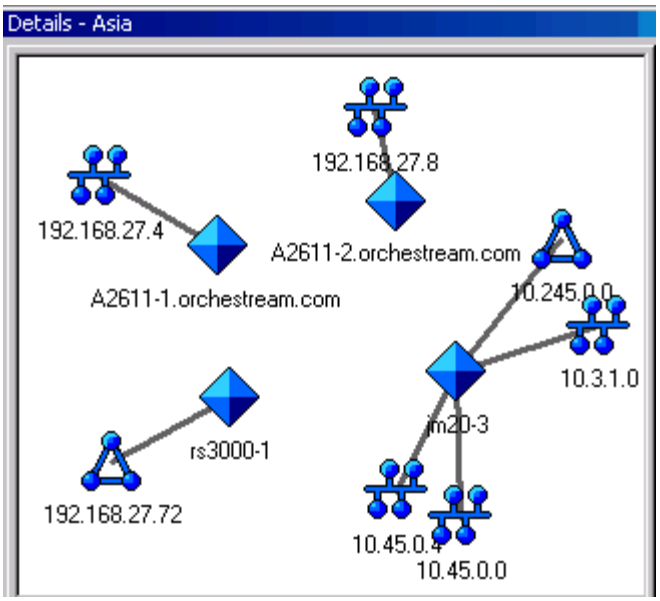

Note: You cannot drag and drop an IPv6 segment when the interface has both an IPv4 and IPv6 address.

VLAN Representation

IP Service Activator discovers VLANs on Cisco Catalyst switches running CatOS and represents the VLAN in the client.

A Catalyst switch may run CatOS (operating at Layer 2) or IOS (operating at Layer 3) or have an MSFC card installed that runs IOS. IP Service Activator supports the discovery of VLANs and port assignment to VLANs on Catalyst switches running CatOS. If a Catalyst switch has an MSFC card that runs IOS, the switch's Layer 2 manageable entity and the MSFC card (Layer 3) are represented as two separate network devices, the Layer 3 entity (IOS), which can be managed by the Cisco device driver, and the Layer 2 entity (CatOS), which can be managed by the CatOS script driver.

VLANs are represented by different icon types on the Layer 2 manageable entity and the MSFC card.

VLAN Representation on a CatOS Device

On a Layer 2 CatOS device, a VLAN is represented as a VLAN interface connected directly to the device. Figure 6–5 shows the VLAN interface icon.

Figure 6–5 The VLAN Interface Icon

The physical ports that are currently assigned to a given VLAN appear as children of the VLAN interface. The property pages for the VLAN are identical to those of an interface object with the addition of a Vlan ID field.

Capabilities are not returned for a VLAN interface on a Layer 2 manageable entity or for the physical ports that participate in the VLAN.

You may apply CDK scripts to the ports that participate in a VLAN by assigning a role to the relevant ports and associating the relevant scripts with that role. You cannot apply CDK scripts to the VLAN interface itself.

VLAN Representation on an MSFC (IOS) Device

On an MSFC (IOS) device, a VLAN is represented as a VLAN interface connected directly to the device. As the VLAN interface represents an IP-addressable Layer 3 entity, a standard interface icon is used to represent the VLAN interface. Figure 6–6 shows the VLAN interface icon for a Layer 3 entity.

Figure 6–6 The VLAN Interface Icon for a Layer 3 Entity

The physical ports associated with the VLAN are not listed as children of the VLAN interface. A VLAN interface on an MSFC card has a network address associated with it and the segment attached to it is therefore displayed as well.

Capabilities are returned for the VLAN interface on an MSFC card, and roles can be assigned in exactly the same way as for physical interfaces.

The VLAN interface on an MSFC card can be associated with a site for inclusion in an MPLS VPN.

Juniper E-series virtual router representation

Juniper E-series virtual routers are displayed on the topology map in a similar way to real devices.

When IP Service Activator discovers a device containing virtual routers it creates a network cloud to 'hold' the routers. Beneath this cloud, IP Service Activator creates device objects that represent both the physical device and its virtual routers. The physical and the virtual routers exist at the same hierarchy level.

You can move virtual routers to other networks in the same way as physical devices.

> **Caution:** Because virtual routers can never be independent of the physical router, you should never move virtual routers on the client, or drag and drop them between networks.

Table 6–1 shows the default naming conventions in IP Service Activator:

Table 6–1 Default Naming Conventions

Items	Naming Convention in Device Driver	Naming Convention in Network Processor
Class-maps r	"ClassMap_Num" Example: class-map match-all "ClassMap_0"	"ClassMap_Num" Example: class-map match-all "ClassMap_0"
Policy-map	"PolicyMap_Num" Example: policy-map "PolicyMap_0"	"PolicyMap_Num" Example: policy-map "PolicyMap_0"
Map-class	"MapClass_Num" Example: map-class frame-relay "MapClass_0"	"MapClass_Num" Example: map-class frame-relay "MapClass_0"

Table 6–1 (Cont.) Default Naming Conventions

Items	Naming Convention in Device Driver	Naming Convention in Network Processor
VC-class	VcClass_Num Example: vc-class atm VcClass_0	VcClass_Num Example: vc-class atm VcClass_0
Acl Name	NamedAcl_Num Example: ip access-list extended NamedAcl_0	NamedAcl_Num Example: ip access-list extended NamedAcl_0
Vrf table Name	Orch_RD Example :ip vrf Orch_1:1876	IPSA_RD Example: ip vrf IPSA_1:1472
RouteMap	RouteMap_Num Example: route-map RouteMap_0 permit 0	SOORoutePolicy_Num Example: route-map SOORoutePolicy_0 permit 0

Creating and Viewing a Topology Map

The map view for a domain is displayed in the Details pane when you double-click on the root network object that represents the entire domain. You can either create the map manually or allow IP Service Activator to populate the map with network objects automatically.

By default, maps are called **Network Map**, but you can change the name to something more descriptive. The map name is displayed on the tab at the bottom of the Details pane. Right-click the tab and select **Properties** to edit the map's properties.

The palette on the right of the screen displays network components that have not yet been added to the map. It can be configured to list all unmapped network components or list components appropriate to the current context. See "Panes in the Domain Management Window" for more information about the palette.

Note: If you are laying out the map manually, you can remove an object from the map by dragging it onto the palette. If you delete an object on the map, it is deleted from IP Service Activator and must be rediscovered.

The Map Toolbar

The map toolbar, shown in Figure 6–7, provides options for working with the topology map.

Figure 6–7 The Map Toolbar

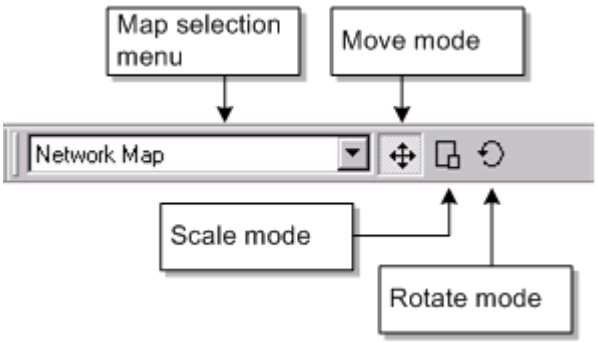

- Map selection menu: lists all the network maps defined within the domain
- Move mode: enables you to move selected nodes to a new location while retaining the layout of the selected nodes
- Scale mode: allows you to enlarge or reduce the area occupied by selected nodes
- Rotate mode: enables you to rotate the selected nodes

See "Selecting and Laying Out Objects on a Map" and "Creating Additional Map Views" for information about using the map toolbar.

Creating the Map Manually

This is the default setting for map creation. In this setup, the map is blank during the discovery process and IP Service Activator adds devices to the palette as they are discovered (the palette is normally displayed to the right of the screen). You can populate the map with network objects by dragging them from the palette or Hierarchy pane on to the map. A snap to grid option is available when mapping objects manually. Connections between objects appear automatically.

Creating the Map Automatically

When automatic layout is selected, IP Service Activator adds network objects to the map as they are discovered.

By default, IP Service Activator adds networks, sites, segments and devices to the map. You can override the default to restrict or expand the network components that are mapped automatically – for example, you can map networks, sites, devices and interfaces automatically. Network components that are not mapped automatically are displayed in the palette.

Note: You cannot drag objects from the palette onto the map while the automatic layout option is selected.

Setting Automatic Layout Parameters

Depending on which object types you select to appear on the map, the number of network objects may be quite large. If the number of objects is too high, this may lead to display problems. For example, objects may overlap and affect the map's clarity or there may be an unacceptable delay time when displaying the map.

IP Service Activator can automatically lay out up to 200 network objects with minimal delay in display time. For greater numbers, there may be some delay in mapping and displaying objects. You can limit the delay by setting the values on the Map Properties dialog box, **Layout** property page including **Max nodes**, **Max fan**, and **Max devices**.

Note: For complete dialog box and property page descriptions, see IP Service Activator online Help.

Note: If the number of network objects exceeds the defined limit, IP Service Activator omits some objects from the map. In this scenario, we recommend you lay out the map manually, incorporating the maximum number of nodes to be included in the map, and use the Recalculate Layout map option to regulate the map's layout. This option is not affected by the Max nodes setting. See "Recalculating a Map Layout" for information about the Recalculate Layout option.

Guidelines for Working with Maps

You might want to work by initially selecting the automatic layout option to create the basic map structure and then change to the manual layout option to add other network objects to the map.

- If you are using a background image with a map, we recommend that you use the manual mapping option so that you can position objects accurately on the image. See "Displaying a Background Image" for information about using a background image.
- If your network includes ATM or Frame Relay VCs, you can drag one VC endpoint on to another to link the endpoints and create a VC link object. Note that you can also link VC objects in the Hierarchy pane. The link between the VC endpoints and the VC link object is shown by a solid connection line.

To specify layout options for a map:

1. With the Map View displayed, right-click anywhere in the Details pane and select **Properties**.

   The Map View dialog box opens.

2. Select the **Layout** property page.
3. If you want to map network objects manually, select **Manually lay out items on map**.

   To specify that objects snap to a grid select the **Snap to grid** option and specify the granularity of the **Grid** in millimeters.

4. If you want to map objects automatically, select **Automatically lay out items on map** and specify the objects to be included on the map including **Networks**, **Sites**, **Devices**, **Interfaces**, **VCs**, **VC End Points**, and **Segments**.
5. Specify values for the map's maximum limits (**Max nodes**, **Max fan**, and **Max devices**).
6. Click **OK**.

To rename a network map:

1. Right-click the map and select **Properties**.

   The Map dialog box opens.

2. Specify the map name.
3. Click **OK**.

Recalculating a Map Layout

If you have the manual layout option selected, the map menu includes an option to recalculate the map layout, that is, arrange objects according to an automatic layout scheme.

Selecting the recalculate option does not switch the automatic layout option on but applies an automatic layout as a one-time operation.

This feature is useful if you have dragged objects on to the map from the palette or moved multiple objects.

To recalculate the map layout:

1. With the Map View displayed in the Details pane and the map's manual layout option selected, right-click the map and select **Recalculate Layout**.

   The Recalculate Layout Options dialog box opens.

2. Select a layout option: **Hierarchical**, **Circular**, or **Symmetric**.
3. Click **OK**.

Selecting and Laying Out Objects on a Map

IP Service Activator provides options for moving, scaling, or rotating a set of selected nodes on a map.

To select objects on a map, do one of the following:

- Click on the network map, hold down the mouse button and drag the cursor over the nodes to be selected.

  As you drag the cursor, a dotted selection box is drawn around the selected nodes.

- Select one or more objects and choose **Select Neighbors** from the objects' context menu.
- Press CTRL while single-clicking each node.

To lay out objects on a map:

1. Select the relevant objects.
2. On the map toolbar, select a mode:
   - If you want to move the node set to a new location on the map, select **Move mode** and drag the nodes to the new location
   - If you want to reduce or enlarge the area occupied by a set of nodes, select **Scale mode** and drag the cursor up (increase area) or down (reduce area)
   - If you want to rotate the selected nodes around a user-selected central node, select **Rotate mode**, click the node that will act as the rotation point, and drag the cursor up (clockwise) or down (anti-clockwise)

Configuring the Palette

By default, the palette is context-sensitive and the objects it lists reflect which type of network object is selected in the map. For example, if you select a device, any unmapped interfaces are displayed in the palette. If you select an interface, the palette lists any unmapped segments, sub-interfaces, and so on. You can turn off this context-sensitive behavior and specify which object types are listed in the palette.

If you drag a displayed object into the palette area, it is removed from the map but not deleted from IP Service Activator.

> **Note:** You can only drag objects between the palette and the map if you have the manual layout option selected.

By default, the palette is always displayed when viewing a map but you can switch off palette display for a map on a temporary or permanent basis.

See "Changing the Size and Position of a Pane" for information on moving and docking the palette.

To specify palette preferences:

1. From the Map menu, select Palette.

   The palette sub-menu is displayed.

2. In the palette sub-menu, do one of the following:

   - Select **Context Sensitive**.

     Makes the palette context-sensitive.

   - Deselect **Context Sensitive** and select the network objects that you want to list in the palette.

To set a permanent palette display option for a map:

1. From the map menu, select **Properties**.

   The Map dialog box opens.

   > **Note:** If a map has a background image, you may need to right-click on the map's tab at the bottom of the Details pane to display its context menu.

2. Click or deselect the **Default Palette Visible** check box as appropriate.
3. Click **OK**.

To set a temporary palette display option for a map:

- With a topology map displayed, from the **View** menu, select **Map Palette**.

The palette is displayed when the **Map Palette** option is checked. Selecting this option shows or hides the palette only for the current map viewing session. If you exit and return to the map, the palette's display is set according to the map's permanent map display setting.

Displaying a Background Image

You can display one or more selected bitmap images as a background to a network topology or VPN map, for example, an outline of a country or region.

If you create your own images, avoid using background colors that are the same as those used by IP Service Activator icons. Background images must be stored in the **ServiceActivator\Backgrounds** directory and must be in Bitmap (.bmp) format.

> **Note:** If there are multiple client components running on a number of host machines, the relevant background image must be accessible to each client component for the image to be displayed

To add or remove a background image:

1. Right-click anywhere within the map area and select **Properties**.

   The Network Map dialog box opens.

2. Select the **Background** property page.

   The **Background Image** list displays all the bitmap files (*.bmp format) that are in the **Backgrounds** folder. If files have already been selected for the map, they are displayed in the list below this field.

3. If you want to add a bitmap file to the map, select the file from the **Background Image** pull-down list and click **Add**.

   The file name is added to the list. Every file in this list will be added to the network map. As you click on each file name, IP Service Activator displays an image **Preview**.

4. If you want to remove a bitmap file from the map, select the file name from the list and click **Remove**.

5. Click **OK** when you are satisfied with the selection.

   The files you selected are added to the network map and positioned at point 0,0 – that is, the top left-hand corner of the map.

> **Note:** If you selected more than one background image, the files are overlaid on the map.

To move an image, do one of the following:

- Click on the image to select it and drag it to a new location.

  A selected image is surrounded by a handled selection box.

- Click on the image to display the Background dialog box and specify the image's new co-ordinates using the **X Axis** and **Y Axis** fields.

> **Note:** The map's 0, 0 co-ordinates are initially at the top left-hand corner of the map. If you move objects beyond this point, they will have a negative co-ordinate value. If you subsequently set a background image's x and y co-ordinates to 0, 0, this may not correspond to the top left-hand corner of the visible map.

To resize a background image by dragging:

1. Select the image.

   A selected image is surrounded by a handled selection box.

2. Click a corner handle on the selection box and drag the handle to resize the image.

To resize a background image using the properties dialog box:

1. From the background image's context menu, select **Properties**.

2. In **Size**, specify the image's width and height.
3. Click **OK**.

Changing the Scale of the Map

The scale of the network map can be varied to enable you to zoom in on a particular section or zoom out to view more of the map.

When a network or VPN map is displayed, a **Map** drop-down menu appears on the menu bar. The menu provides temporary zoom values for the map, enabling you to change the scale of the map between 25% and 200%. Alternatively, you can select the auto zoom feature to scale the map according to the value of its length or width, whichever is the greatest value, and retain the aspect ratio.

The temporary zoom setting is not retained if you exit and review the map. To make the zoom setting more permanent, set the map's default zoom setting. By default, this is 100%.

To set a map's temporary zoom setting, from the **Map** menu, select **Zoom**, and then the magnification that you want.

You can change the scale of the map between 25% and 200%. Selecting **Auto Zoom** scales the map according to its length or width, whichever is the greatest value.

To define a map's default zoom setting:

1. Right-click the map and select **Properties**.

   The Map dialog box opens.

   Note: If a map has a background image, you may need to click on the map's tab at the bottom of the Details pane to display its context menu.

2. In the **Default Zoom** field, select a new zoom value.
3. Click **OK**.

   The map is resized to the new zoom setting.

Creating Subsidiary Networks and Maps

To simplify the management of large and complex networks you can divide the network into a number of subsidiary networks, each with its own topology map. Depending on the complexity of your network and the information you want to represent, you can create a multi-level hierarchy of subsidiary networks. Each device within the domain can only appear in one network.

To create a subsidiary network

1. Do one of the following:
   - If you want to create a new network map that initially contains no devices, right-click on a blank area in the map display window for the domain and select **Add Network**.
   - If you want to create a new network map that contains a copy of part of an existing network map, click and drag over the relevant devices to select them and select **Add Network**.

The Network dialog box opens.

2. Enter a name for the subsidiary network and a brief description of the network and click **OK**. You can enter the rest of the network properties after you have created the map and decided what devices are to be included in the network.

A network cloud object appears on the map representing the subsidiary network. The network cloud object also appears in the Hierarchy pane.

> **Note:** Alternatively, you can select the top-level network in the Hierarchy pane and select Add Network from the context menu. In this case the network cloud will appear in the palette but will not initially appear on the map.

To move devices between networks:

- To move a device into a lower-level network, select the device and drag it onto the appropriate network cloud. When you drop the device icon, it disappears into the new network cloud leaving only its links to other network objects visible.
- To move a mapped device into a higher-level network, right-click the device in the Details pane and select **Promote Upward**. It is removed from the current map and appears in the higher-level map's palette.

  Any links you made between devices in the lower-level network are not retained.
- Alternatively, you can move devices between networks by dragging and dropping devices within the Hierarchy pane.

To set up and view maps for subsidiary networks:

A map is automatically created for each subsidiary network.

- To view a subsidiary map, double-click on the new network cloud (either on the Hierarchy pane or on an existing map).

  Initially subsidiary maps are blank, and you need to create the map by dragging devices from the palette.
- To return to a higher level map, click the **Up** icon.

Creating Additional Map Views

In addition to creating subsidiary networks and maps, you can set up several map views of each network. These appear as separate tabs on the Details pane when the map is displayed. For example, you can create one map that just displays the network devices and segments and another that also displays device interfaces and sub-interfaces. Devices and other network objects can appear on more than one map view of the network.

To create an additional network map for the domain:

1. Do one of the following:
   - If you want to create a new network map that initially contains no devices, right-click a blank area on the map and select **Add Map View**.
   - If you want to create a new network map that contains selected devices, click and drag over the devices to select them, and right-click and select **Add Map View**.

   The Map View dialog box opens.

2. Enter a name for the map and a brief description of the map view and click **OK**. A new tab appears at the bottom of the Map display window.
3. Click on the tab to display the new map view. To start with, the map area is blank and all the network objects that have been discovered in the domain are listed in the palette.
4. Create the map in the standard way by dragging devices from the palette.

To view another network map, on the map toolbar, select a map from the network map selection menu.

7

Checking Device Status and Capabilities

This chapter describes how to:

- View the status of a device or interface
- View the list of interfaces on a device
- Check a device or interface's capabilities
- Define device/interface capabilities for devices configured by a Network Processor cartridge

Viewing the Status of a Device or Interface

While a device is managed by Oracle Communications IP Service Activator, the proxy agent polls it regularly to check that it is still running. A device's status, as reported in the IP Service Activator client, may change as a result of polling. For example, if a device fails, its status is changed to Unreachable. The device is still polled and will be reset when it recovers. If the problem fails to reset, however, or if a critical fault is raised on a device, the status is changed to Intervention Required until the problem is resolved.

Figure 7–1 shows the range of device states and the transitions between them.

Note that initial discovery of a device may be by any of the following routes:

- Discovering or creating the device through the client. See "Running Device Discovery" for information.
- Discovering or creating the device via the OSS Integration Manager (OIM) – see *IP Service Activator OSS Integration Manager Guide*.

You can check the status of the devices and interfaces within a domain by viewing the topology map in status context. Every device and interface state is represented by a display color in the map.

See "How Objects are Represented" for more information.

Figure 7–1 Device States and Transitions

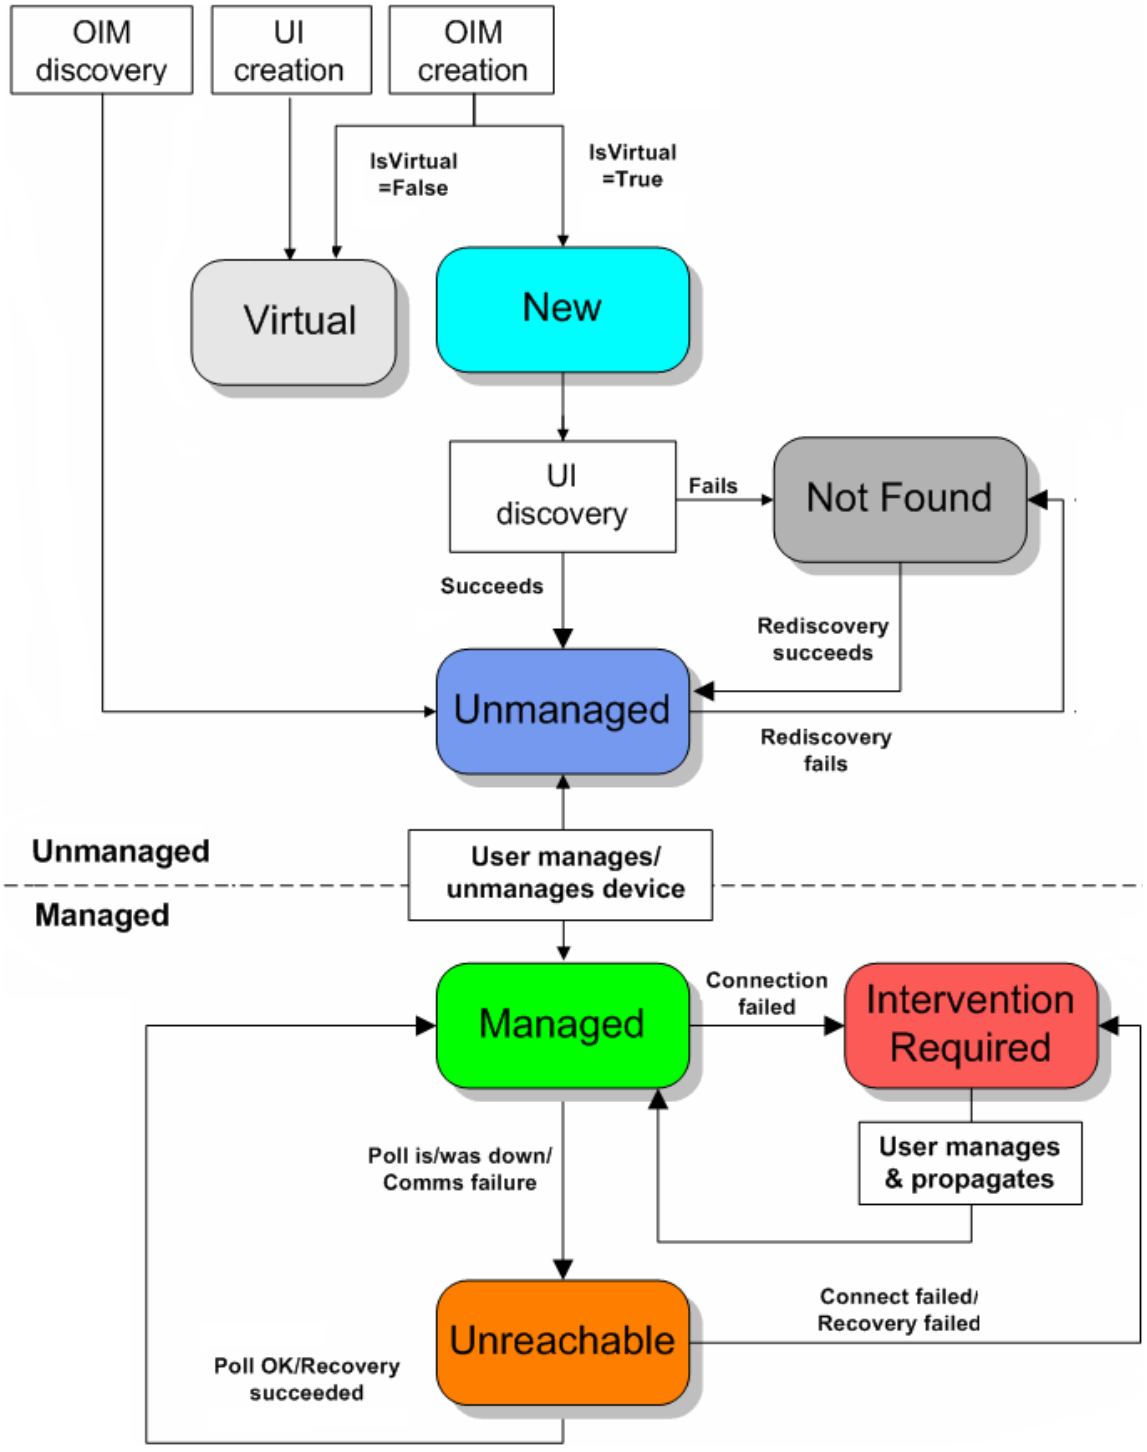

To access a device or interface's properties dialog box:

1. Right-click the object, either on the network map or **Topology** tab, and select **Properties**.

If the selected object is a device or an interface, the information displayed in the Capabilities property page is particularly important. It summarizes the support that the client provides for QoS mechanisms, VPN capabilities, policy rules and SLA

monitoring. You should be aware of these capabilities when planning services and policies. See "Checking Capabilities" for more information about interface capabilities.

Viewing a List of Interfaces on a Device

You can list the interfaces that are available on a device by double-clicking the device in the Hierarchy pane. IP Service Activator lists the selected device's interfaces in the Details pane.

You can also list a device's interfaces on the Interfaces property page on the Device dialog box.

For each interface on the device, the Interfaces property page displays the information in Table 7–1.

Table 7–1 Interfaces Property page

Information Displayed	Description
Description	A description of the connection. This is the SNMP ifDescription parameter.
Number	The number of the interface.
IP Address	IP address of the interface. (If the interface has multiple IP addresses, this is the first address.) The IP Address field accommodates both IPv4 and IPv6 addresses.
State	Status of the interface (Up, Down, Shutdown or Unknown). For virtual interfaces, the state is always Unknown.
Type	The type of interface, such as Ethernet, point-to-point. This is the SNMP ifType parameter.
Role As	The policy role assigned to this interface, such as Core, Local, Access, None or Disabled.

Note: Double-clicking on an interface in the Description column opens the properties dialog box for that interface.

Checking Capabilities

IP Service Activator displays information about the capabilities of devices, interfaces, sub-interfaces, and VC endpoints. Being aware of each device and interface's capabilities is an important factor in developing and applying policies and services because they indicate the VPN service, policy, and measurement types that are supported. In addition, if capabilities are not known, rules, PHB groups, and some measurement types cannot be applied to that network component.

Capabilities are obtained as part of the device discovery process. However, if IP Service Activator is unable to obtain the capabilities, for example because the device's security settings are incorrect, you can request the capabilities later.

IP Service Activator derives interface and sub-interface capabilities from the device's operating system, device type, and interface name. VC endpoint capabilities also take account of the interface type.

IP Service Activator stores capability details for devices, interfaces, sub-interfaces, and VC endpoints.

Device-level Capability Categories

A device may have any of the following capabilities:

- SAA support: indicates the maximum number of SAA probes (operations) supported and which SAA operations are supported – ICMP Echo, UDP Echo, Jitter or TCP Connect
- NetFlow support: indicates in which versions of UDP format NetFlow data may be exported from the device – Version 1, Version 5 or Version 8

These capabilities are listed on the Capabilities property page of the device's properties dialog box.

If the words 'Capabilities not obtained' are displayed on this page, IP Service Activator has not obtained the relevant information from the device. This may be because IP Service Activator could not communicate with the device, because the security parameters are incorrect (for Cisco devices) or because the device's operating system is not supported.

Interface-level Capability Categories

An interface, sub-interface, or VC endpoint may have any of the following capabilities:

- Access: indicates that you can implement access rules.
- ATM: indicates that you can implement PHB groups using the ATM mechanism.
- CCC: indicates support for CCCs (Circuit Cross Connects).
- Class-based queuing: indicates support for CB-WFQ.
- Classification: indicates that you can implement classification rules.
- FRTS: indicates that you can implement PHB groups using the Frame Relay Traffic Shaping mechanism.
- Martini: indicates support for Layer 2 Martini VPNs
- MQC: indicates support for Modular QoS CLI on Cisco devices.
- Policing: indicates that you can implement policing rules.
- Priority Queuing: indicates that you can implement PHB groups using the Priority Queuing mechanism.
- Rate Limiting: indicates that you can implement PHB groups using rate limiting.
- Virtual Circuits: indicates a Frame Relay or ATM VC.
- VPN: indicates that this interface supports VPNs.
- WFQ: indicates that you can implement PHB groups using the WFQ queuing mechanism.
- WRED: indicates that you can implement PHB groups using the WRED queuing mechanism.
- WRR: indicates that you can implement PHB groups using the WRR queuing mechanism.

These capabilities are listed on the Capabilities property page of the Interface dialog box.

If the words 'Capabilities not obtained' are displayed on this page, IP Service Activator has not obtained the relevant information from the device. This may be because IP Service Activator could not communicate with the device, because the security

parameters are incorrect (for Cisco devices) or because the device's operating system is not supported.

If no text is displayed in a field, it indicates that the interface does not support any IP Service Activator features or the device or interface are not supported. If the device or interface are not supported, an error message is displayed in the current faults pane.

You can find out more detailed information about most of the capabilities listed on the Capabilities property page by double-clicking on the icon next to the capability. The details are displayed beneath the capability.

The details displayed vary depending on the capability type. The ways in which the interface can identify and classify traffic is specified for rule and queuing capabilities, while for queuing mechanisms, the number of queues that can be handled by the interface is specified.

Refetching Device Capabilities

If IP Service Activator is unable to discover capabilities during the discovery process – for example, because the device's security settings are incorrect – you can request the capabilities at a later point.

In addition, if the device capabilities have changed – for example, as a result of an operating system upgrade – you can refresh capabilities by resetting and remanaging the device.

Note: Changes made to the capabilities files do not take effect in the system until the fast start process has completed. Erratic behavior might occur if there is any attempt to re-fetch capabilities while the proxy agent and/or device driver is in the start-up process. Wait until the message Driver is up and running appears in the Current Faults Pane before rediscovering capabilities.

Although the network processor does not include a fast-start process, you must wait for startup to complete before re-fetching capabilities. Network processor start is indicated in an information message in the **Current Faults** Pane.

To discover unknown device capabilities:

1. From the Discover menu, select Discovery.

   The Discovery dialog box opens.

2. Select Fetch Capabilities from the Discovery Type drop-down.

3. Click OK or Apply.

IP Service Activator attempts to discover capabilities for all devices that have not had capabilities returned. It does not rediscover capabilities where they have already been reported.

To update previously found device capabilities:

1. Unmanage the device ensuring IP Service Activator configuration is left on the device:

   a. Open the device's property pages and select the Management property page.

   b. Ensure that Inherit System Action is set to Leave.

c. Click the Unmanage button.

d. Click OK to close the dialog box.

e. Commit the transaction.

2. Reset the device capabilities:

a. Open the device's properties dialog box and select the Discovery property page.

b. Click on the Reset Device Capabilities button.

c. Click OK.

d. Commit the transaction.

3. Fetch the updated capabilities:

a. From the Discovery menu, select Discover.

The Topology Discovery dialog box opens.

b. In the Discovery Type field, select Fetch Capabilities.

c. Click OK.

d. Commit the transaction.

4. Remanage the device:

a. Open the device's properties dialog box and select the Management property page.

b. Click the Manage button.

c. Click OK.

d. Commit the transaction.

Modifying Device/Interface Capabilities

You can modify the interface capabilities of devices configured by a specific cartridge unit by modifying the capabilities file registered to that cartridge unit. For more information, see the cartridge guide that applies to your IP Service Activator installation.

> **Note:** You can modify the device and interface capabilities only for devices configured by a network processor cartridge.

The **MIPSA_registry.xml** file lists information about the cartridge units in a vendor-specific cartridge family. It names the cartridge unit name, driver type, device type, OS type, some service model/device model/CLI transforms, and the applicable capabilities file. A sample entry follows:

```
<!-- Cisco 3640 devices with IOS 12.2(8)T4 -->
<cartridgeUnit>
    <name>com.metasolv.serviceactivator.cartridges.cisco.units.cu1.3640.
        12.2(8)T4</name>
    <driverType>cisco</driverType>
    <deviceType>Cisco 3640</deviceType>
    <osRegex>.*12\.2\(8\)T4.*</osRegex>
    <smToDmQuery>com/metasolv/serviceactivator/cartridges/cisco/units/cu1/
        sm2dm.xq</smToDmQuery>
```

```
    <dmValidation>com/metasolv/serviceactivator/cartridges/cisco/units/cu1/
         dmValidation.xq</dmValidation>
    <dmToCliQuery>com/metasolv/serviceactivator/cartridges/cisco/units/cu1/
         annotatedDm2Cli.xq</dmToCliQuery>
    <capabilities>com/metasolv/serviceactivator/cartridges/cisco/capabilities/
         cisco_3640_12.xml</capabilities>
</cartridgeUnit>
```

In the above sample, **MIPSA_registry.xml** defines a cartridge unit used to configure Cisco 3640 devices running the 12.2(8)T4 IOS. It names the capabilities file **cisco_3640_12.xml**.

Each capabilities file consists of sections, including the following:

- device capabilities
- interface capabilities, inbound and outbound sections
- sub-interface capabilities, inbound and outbound sections
- virtual circuit capabilities, inbound and outbound sections

In file **cisco_3640_12.xml**, for example, you can reference the enumerated types listed in Table 7–2 to add Classification and Marking capabilities on interfaces:

Table 7–2 Enumerated Types Referenced in cisco_3640_12.xml

Enumerated Types	Description
SRC_IP	Source IP address
DST_IP	Destination IP address
SRC_PORT	Source port
DST_PORT	Destination port
IP_PROTO	Internet Protocol
DIFFSERV	Differentiated Services
IPv4PRECEDENCE	IP version 4 precedence bits
IPv4TOS	IP version 4 type-of-service byte (octet)
URL	Universal Resource Locator
MIME	Multipurpose Internet Mail Extensions
APP_PROTO	Application protocol
APP_NAME	Application name
DOMAIN_NAME	Domain name
IEEE_802_1_PBITS	IEEE 802.1P priority field
MPLS_EXP	Multi-Protocol Label Switching experimental value
MPLS_EXP	Multi-Protocol Label Switching experimental value
FR_DE	Frame Relay discard eligibility bit
ATM_CLP	Asynchronous Transfer Mode cell loss priority bit
CLASS_MAP	Class map
MATCH_ANY	Match any
ALCATEL_INTERNAL_QUEUE	Alcatel internal queue
TCP_HEADER_OPTIONS	Transmission Control Protocol header options

Table 7–2 (Cont.) Enumerated Types Referenced in cisco_3640_12.xml

Enumerated Types	Description
EXCLUDE_CLASSIFICATION	Exclude classification

Code Sample

Sample code excerpts follow. (Entries removed are implied by three stacked periods, for the sake of brevity.)

```
<caps:capabilities
xmlns:caps="http://www.metasolv.com/serviceactivator/capabilities"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <caps:device>
  ...
    <caps:ipsec_support>
      <caps:esp_algorithms_supported>0</caps:esp_algorithms_supported>
      <caps:ah_algorithms_supported>0</caps:ah_algorithms_supported>
      <caps:compression_algorithms_supported>0</caps:compression_algorithms_
supported>
      <caps:ipsec_modes_supported>0</caps:ipsec_modes_supported>
    </caps:ipsec_support>
    ...
  </caps:device>
  <caps:interface>
    <caps:ifType>32</caps:ifType>
    <caps:inbound>
      ...
      <caps:service_rule_support>
        <caps:rules_supported>0</caps:rules_supported>
        <caps:mark_codepoints_supported/>
        <caps:limits_supported>0</caps:limits_supported>
        <caps:guarantees_supported>0</caps:guarantees_supported>
        <caps:limits_and_guarantees_supported>0</caps:limits_and_guarantees_
supported>
        <caps:classification_supported/>
      </caps:service_rule_support>
      <caps:policing_rule_support>
        <caps:policing_supported>7</caps:policing_supported>
        <caps:classification_supported>
          <caps:supported>SRC_IP</caps:supported>
          <caps:supported>DST_IP</caps:supported>
          <caps:supported>SRC_PORT</caps:supported>
          <caps:supported>DST_PORT</caps:supported>
          <caps:supported>IP_PROTO</caps:supported>
          <caps:supported>APP_PROTO</caps:supported>
        </caps:classification_supported>
        <caps:classification_queues_supported>7</caps:classification_queues_
supported>
        <caps:mark_codepoints_supported>
          <caps:supported>SRC_IP</caps:supported>
          <caps:supported>DST_IP</caps:supported>
          <caps:supported>SRC_PORT</caps:supported>
        </caps:mark_codepoints_supported>
      </caps:policing_rule_support>
      ...
    </caps:inbound>
    <caps:outbound>
      ...
      <caps:service_rule_support>
```

```
        <caps:rules_supported>0</caps:rules_supported>
        <caps:mark_codepoints_supported/>
        <caps:limits_supported>0</caps:limits_supported>
        <caps:guarantees_supported>0</caps:guarantees_supported>
        <caps:limits_and_guarantees_supported>0</caps:limits_and_guarantees_
supported>
        <caps:classification_supported/>
      </caps:service_rule_support>
      <caps:policing_rule_support>
        <caps:policing_supported>7</caps:policing_supported>
        <caps:classification_supported>
          <caps:supported>SRC_IP</caps:supported>
          <caps:supported>DST_IP</caps:supported>
          <caps:supported>SRC_PORT</caps:supported>
          <caps:supported>DST_PORT</caps:supported>
          <caps:supported>IP_PROTO</caps:supported>
          <caps:supported>APP_PROTO</caps:supported>
        </caps:classification_supported>
        <caps:classification_queues_supported>7</caps:classification_queues_
supported>
        <caps:mark_codepoints_supported>
          <caps:supported>SRC_IP</caps:supported>
          <caps:supported>DST_IP</caps:supported>
          <caps:supported>SRC_PORT</caps:supported>
        </caps:mark_codepoints_supported>
      </caps:policing_rule_support>
      ...
    </caps:outbound>
    <caps:subInterface>
      <caps:inbound>
      ...
      </caps:inbound>
      <caps:outbound>
      ...
      </caps:outbound>
    </caps:subInterface>
  </caps:interface>
</caps:capabilities>
```

8

Working with Transactions

This chapter describes how to work with transactions in Oracle Communications IP Service Activator. It explains how to create, commit, save and schedule, merge and unmerge, and delete transactions.

For an overview of transactions, see *IP Service Activator Concepts*.

Creating Transactions: The Current Transaction

When you work in the IP Service Activator client, the changes that you make to the object model are held in the current transaction. This transaction acts as a cache, queuing the changes that you make. The options available for stopping the transaction can include some or all of the following:

- Saving the transaction: details of the transaction are added to the database but the transaction's changes to the object model and device configuration are not implemented
- Committing the transaction: the common object model is updated with the transaction's changes and any associated configuration is propagated to the network
- Scheduling the transaction: similar to saving the transaction, but a date and time are set for implementing the transaction's changes
- Discarding the transaction: the changes held in the current transaction are discarded

Some or all of these options might be available depending on your security rights. For example, some users can save only the current transaction, whereas users with a higher security level can commit or schedule transactions. For information about defining user security levels, see *IP Service Activator System Administrator's Guide*.

Figure 8–1 illustrates these options and how they are queued in the current transaction.

Figure 8–1 Transaction End Options

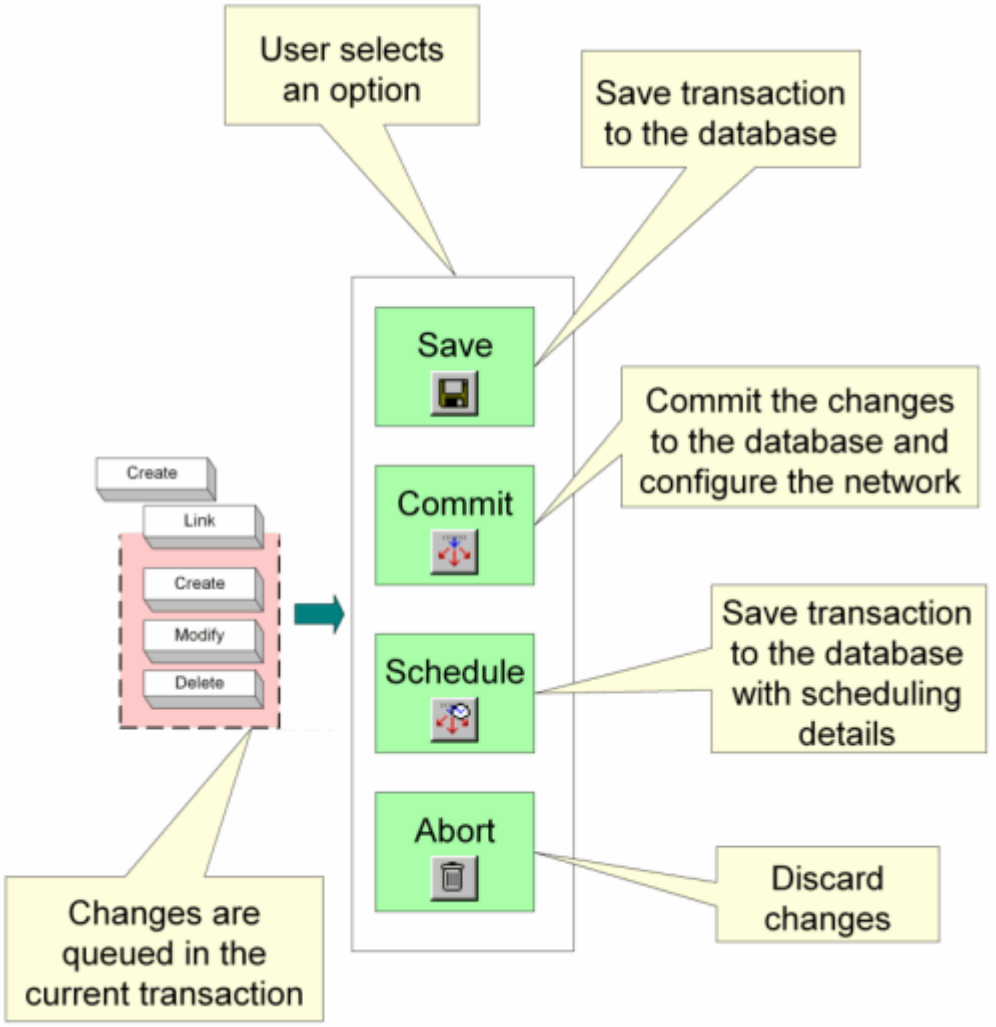

Note: You can undo commands during the current transaction up to the start of the transaction.

Network discovery and related tasks, such as fetching capabilities, cannot be done part way through the current transaction. These tasks can be done only immediately after a transaction is saved or committed and before a new transaction is started.

Committing Transactions

When you commit a transaction, IP Service Activator validates the transaction changes. The Transaction dialog box opens, showing details of the transaction, including any concrete objects or faults that will be generated by confirming the commit.

After reviewing the potential concrete objects and faults, you can choose whether to proceed with the commit or cancel.

If you proceed with the commit, IP Service Activator updates the common object model with the transaction's changes, and configuration changes are propagated to the network. If there are other users working on remote user interfaces, IP Service

Activator updates the object model on each remote client according to the updated common object model. For information about local and common object models, see *IP Service Activator Concepts*.

If you cancel the commit, you can revise details of the policy or the points at which it is applied before trying to commit the transaction again.

Note: Before you can commit a saved transaction, you must merge it. See "Merging or Previewing Transactions"for more information about merging.

If a user on a remote machine is creating a transaction that conflicts with the updated object model, a warning message is displayed and unsaved changes made by that user are discarded.

When you commit a transaction, IP Service Activator creates a new transaction object in the **Committed Transactions** folder. The transaction bundles the actions that have been committed.

A single committed transaction object might correspond to a number of previously saved and merged transactions. For example, if you select a number of transactions, merge them, and commit them, only one committed transaction object is created for those transactions.

Figure 8–2 shows the **Transactions** folder with three pending transactions. Figure 8–3 shows the Committed Transactions folder with a single committed transaction object that represents the three merged transactions.

Figure 8–2 Pending Transactions

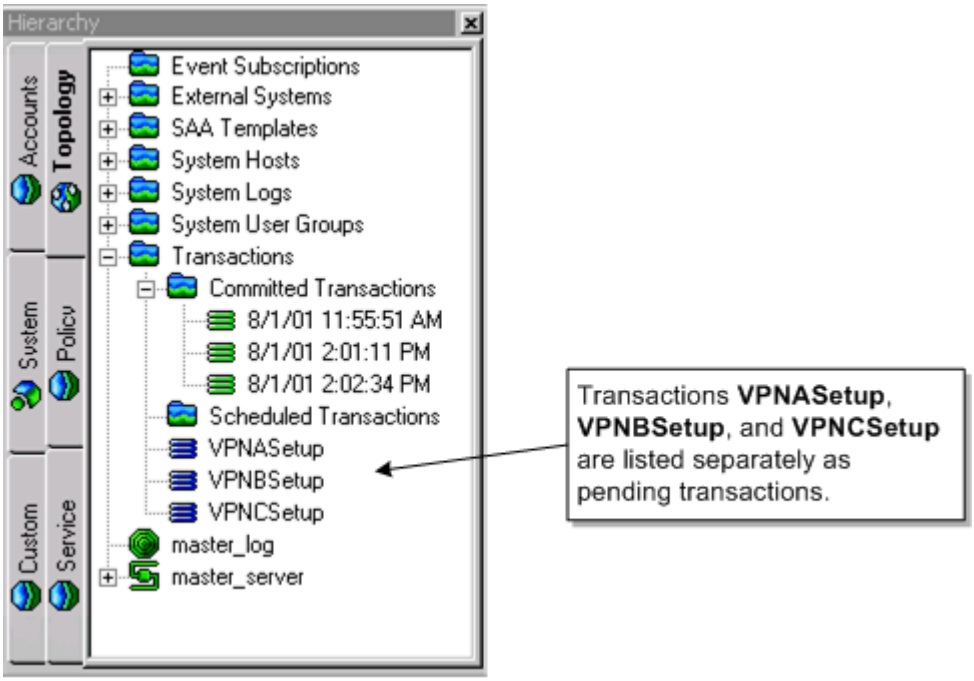

Figure 8–3 Merged and Committed Transactions

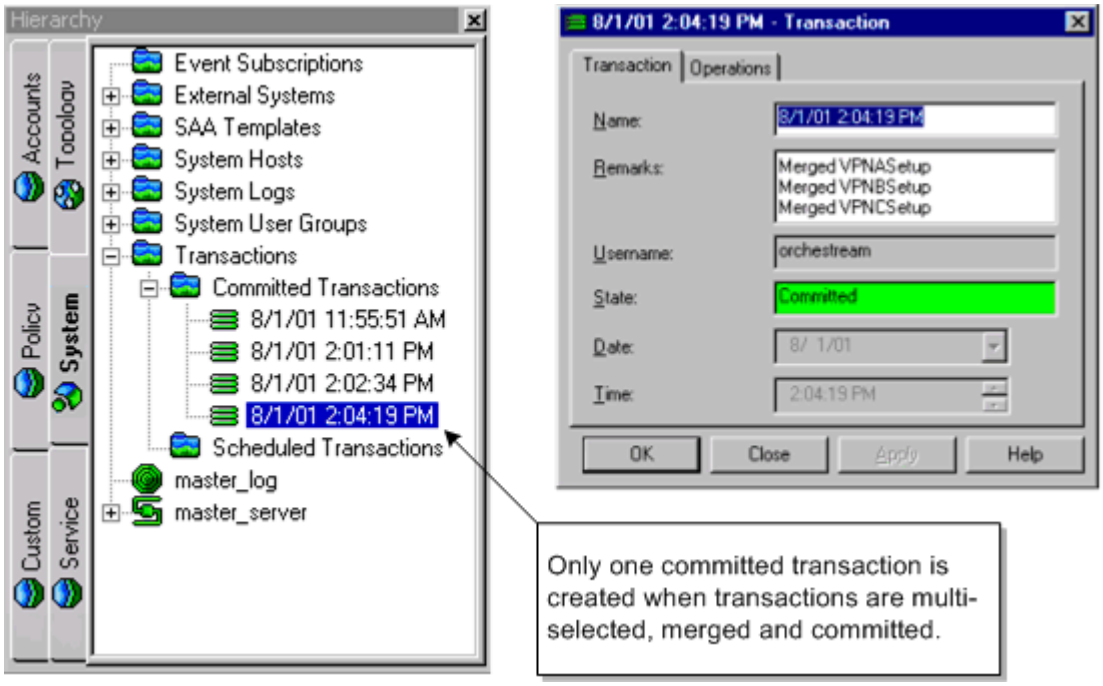

After you commit a transaction, the transaction buttons on the toolbar are disabled. When the next change is made in the client, IP Service Activator starts a new current transaction and the toolbar buttons are enabled.

If necessary, you can remove a committed transaction's changes from the object model and any related configuration from the network by unmerging the transaction. See "Unmerging or Rolling Back Transactions" for more information.

To commit a transaction:

1. From the toolbar, click the **Commit** icon.

   The **Commit** icon is enabled only if you have unsaved changes in IP Service Activator, and your user security level permits you to commit transactions.

   The Transaction dialog box opens.

   IP Service Activator applies a default name that indicates the date and time at which the transaction was committed and populates the **Remarks** field with information about the committed transaction, listing and naming any merged transactions. Changes that are saved as a transaction before being committed are listed as 'unidentified user action(s)'.

2. (Optional) Select the **Concretes** property page and review the concrete objects that will be created if you proceed with the commit.

   Double-clicking on a listed concrete object displays details of the relevant policy target in the Details pane.

3. (Optional) Select the **Fault** property page and review the faults that would result from committing the transaction.

4. If you want to cancel the commit as a result of your review, click **Cancel**.

5. (Optional) Change the transaction's default details on the **Transaction** property page based on the review of concretes and faults.

6. Click **OK**.

   The transaction is saved to the database, the common object model is updated and any related configuration changes are propagated to the network. The common object model is pushed to remote user interfaces.

Saving Transactions

Saving a transaction suspends the transaction changes to a future date or time.

When you save the current transaction, IP Service Activator creates a transaction object and adds it to the common object model's transaction store. The object model changes associated with that transaction are not made and no configuration changes are propagated to the network.

For the user who saves or schedules the current transaction, the client shows the state of the network as it is currently configured, that is, the state of the system held by the common object model.

For example, five sites are created through the client and the current transaction is then saved and held in a pending state.

Figure 8–4 shows the Hierarchy pane before the transaction is started.

Figure 8–4 The Hierarchy Pane Before the Transaction

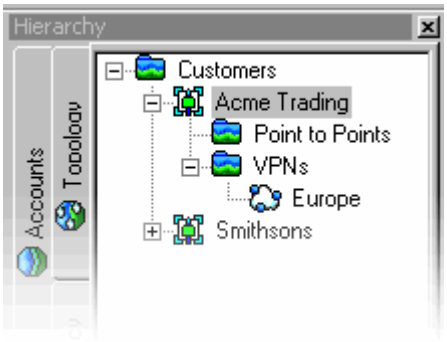

Figure 8–5 shows the Hierarchy pane after starting a new transaction and creating five sites: Amsterdam, Berlin, Madrid, Paris, and Stockholm. The client shows the changes made in the current transaction.

Figure 8–5 The Hierarchy Pane After Starting a New Transaction

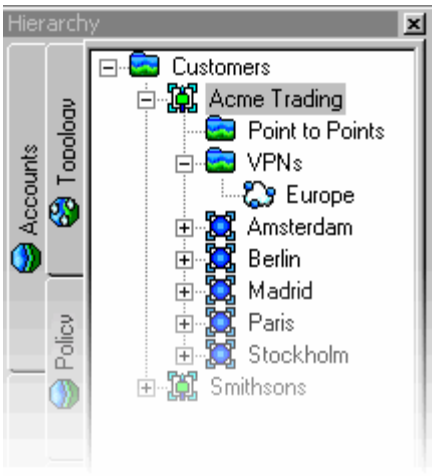

After saving the transaction, the client reflects the network as it is currently configured. In the case of this example, the client reverts to Figure 8–4.

When you save the current transaction, its status changes to Pending.

After saving a transaction you can:

- Merge it into the current model: this is an essential step before committing the transaction. See "Merging or Previewing Transactions" for information about merging.
- Schedule the transaction. See "Scheduling Transactions" for information.
- Delete the transaction. See "Deleting Transactions" for information.

To save a transaction:

1. On the toolbar, click the **Save** icon.

   The **Save** icon is enabled only if there are unsaved changes in the client. The Transaction dialog box opens.

2. Enter the name and remarks.
3. Click **OK**.

   The transaction is saved to the common object model's transaction store. The client reverts to show the state of the network as it is currently configured, that is, the state of the system held by the common object model.

Note: If a transaction incorporates a merged transaction — for example, a pending transaction was merged during the current transaction — icons representing both the merged transaction and the saved transaction are displayed on the System tab when the 'parent' transaction is saved.

Scheduling Transactions

Like saved transactions, a scheduled transaction's changes are put on hold but you can set the date and time for implementing the transaction's changes. Depending on your user security level, you may be able to:

- Schedule the current transaction

- Schedule a pending, merged, or unmerged transaction

When you schedule a transaction, its status changes to Scheduled.

If a scheduled transaction runs successfully, IP Service Activator updates the common object model and propagates the changes to the network. A new committed transaction object is added to the transaction store and is displayed in the **Committed Transactions** folder.

If a scheduled transaction fails, the transaction is moved to the list of pending transactions and its status changes to Failed.

After scheduling a transaction you can:

- Leave the transaction to be committed at the allocated date and time
- Merge and commit the transaction manually. See "Merging or Previewing Transactions" and "Committing Transactions" for information.
- Unschedule the transaction. See "Scheduling Transactions" for information.

To schedule the current transaction:

1. From the toolbar, select the **Schedule** icon.

   The **Schedule** icon is enabled only if you have started a transaction. The Transaction dialog box opens.

2. Enter the details, including name, remarks, date, and time.
3. Click **OK**.

   The current transaction is saved to the database. The client reverts to show the state of the network as it is currently configured – that is, the state of the system held by the common object model.

To schedule a pending transaction:

1. On the **System** tab, open the **Transactions** folder.
2. Right-click a pending transaction and select **Schedule**.

   The Transaction dialog box opens.

3. Enter a date and time to specify when to commit the transaction.
4. Click **OK**.

   IP Service Activator moves the transaction to the **Scheduled Transactions** folder.

To unschedule a transaction:

1. On the **System** tab, open the **Transactions** folder and select a scheduled transaction.
2. From the transaction's context menu, select **Unschedule**.

   IP Service Activator moves the transaction out of the **Scheduled Transactions** folder and changes its status to Pending.

Note: You must select the **Commit** button after scheduling or unscheduling a transaction in order to implement the scheduling details.

Merging or Previewing Transactions

If a transaction is saved or scheduled, you can preview its changes before committing the transaction.

> **Note:** For saved transactions, you must merge the transaction before it can be committed.

Previewing a transaction merges its changes into the local object model, that is, the object model maintained by the local client. The merge does not affect the common object model, and the results are therefore not reflected in remote user interfaces. This concept is illustrated in Figure 8–6. For information about local and common object models, see *IP Service Activator Concepts*.

Figure 8–6 Transaction Merged in Local Object Model but not in Common Object Model

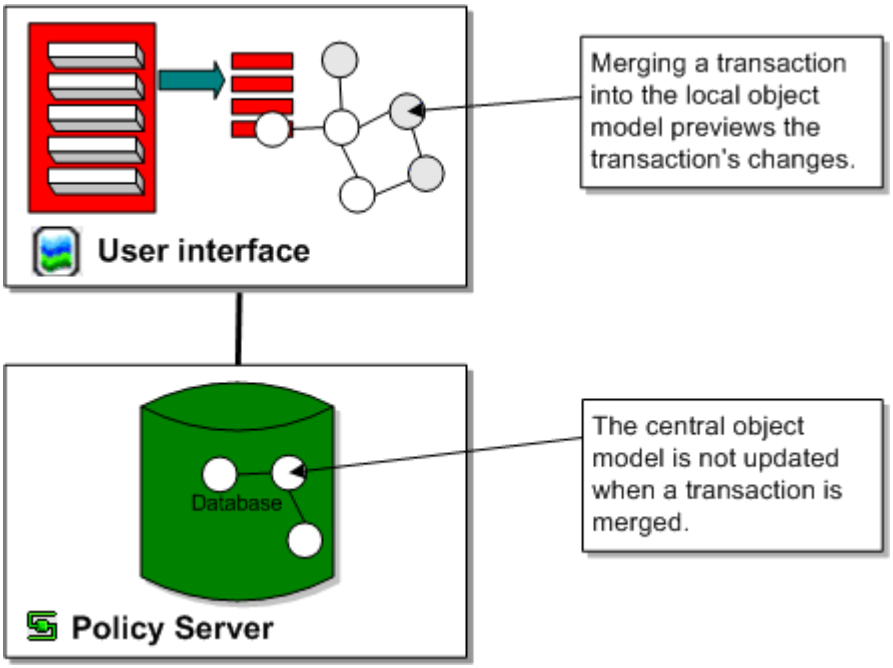

During the merge, IP Service Activator tests the validity of the transaction against the local object model. If there is a conflict, IP Service Activator abandons the merge and reports an error; this occurs, for example, if you attempt to merge a transaction that creates an object that already exists.

You can merge pending or scheduled transactions. If you merge a scheduled transaction, IP Service Activator removes its associated schedule details. The transaction can be committed manually or rescheduled if required.

You can multi-select a number of transactions and merge them in one step. IP Service Activator merges the transactions in the order in which they were selected. See "Selecting Transactions" for information.

After merging a transaction you can do one of the following:

- Commit the merged transaction's configuration changes. See "Committing Transactions" for more information.

- Perform additional tasks and save or commit the new transaction. See "Saving Transactions" and "Committing Transactions" for more information.
- Abort the current transaction: this unmerges the merged transaction and removes its changes from the local object model. See "Discarding the Current Transaction" for more information.

To merge a transaction:

1. On the **System** tab, open the **Transactions** folder.
2. Right-click a pending transaction and select **Merge Changes**.

   The changes associated with the transaction are merged into the local object model. The merged transaction's icon is no longer displayed on the **System** tab.

Unmerging or Rolling Back Transactions

When you roll back a transaction, its changes are removed from the object model and, where configuration has been installed on network devices, it is removed.

The ability to roll back a committed transaction depends on whether subsequent changes are dependent on that transaction's changes. For example, a transaction that creates a VPN to which sites and interfaces have subsequently been linked cannot be rolled back. In this case, the client displays an error message.

Rolling back a transaction is a two-step process that consists of:

1. Unmerging the transaction from the local object model.
2. Committing the unmerge to update the common object model and remove any associated configuration details from the network. See "Committing Transactions" for information about committing transactions.

To unmerge a transaction:

1. On the **System** tab, open the **Transactions** folder.
2. Right-click a committed transaction, and select **Unmerge Changes**.

   IP Service Activator attempts to remove the transaction's changes from the local object model. A message indicates whether the action was successful.

Discarding the Current Transaction

Aborting a transaction stops the current transaction and discards its changes.

To abort the current transaction:

1. On the toolbar, click the **Abort** icon.

   IP Service Activator asks you whether you want to lose the changes.
2. Click **Yes**.

The current transaction is stopped and its changes discarded. The client reverts to show the state of the network as it is currently configured, that is, the state of the system held by the common object model.

Deleting Transactions

You can delete a Pending, Failed, or Scheduled transaction.

> **Note:** If you want to undo the changes made by a transaction, unmerge the transaction and commit the change. See "Unmerging or Rolling Back Transactions" for information.

To delete a transaction:

1. On the **System** tab, open the **Transactions** folder.
2. Right-click a transaction, and select **Delete**.
3. Click the **Commit** icon.

   The deletion completes.

Managing Transactions

This section describes broad management tasks associated with transactions. It describes how to:

- Run a client in confirmed commit mode
- View a list of transactions
- Search committed transactions
- View transaction details
- Check which user committed a transaction
- Specify the transaction archive limit

Running in Confirmed Transaction Mode

By default, a transaction that has been committed is queued and processed by the policy server as a background task. This means that a user can continue working in the client and commit further transactions while the policy server saves data to the database. However, if the policy server fails, transactions that it has not written to the database are lost.

In confirmed transaction mode, a user who has committed a transaction cannot make any further changes in the client until the transaction is successfully saved to the database.

> **Note:** If your deployment integrates IP Service Activator with other OSS systems, it is particularly important that these systems are guaranteed persistence of data. This ensures synchronization between systems. You may therefore want to consider running IP Service Activator user interfaces in confirmed commit mode in this scenario.

To run a client in confirmed transaction mode:

1. From the **Tools** menu, select **Options**.

   The Options dialog box opens.
2. Select the **Transactions** property page and select the **Confirmed transaction mode** check box.
3. Click **OK**.

Searching Committed Transactions

IP Service Activator supports the ability to search committed transactions.

To search committed transactions directly from the Hierarchy pane:

1. In the Hierarchy pane, click the **System** tab.
2. Expand the **Transactions** folder.
3. Right-click the **Committed Transactions** folder, and select **Find From Here**.

   The Find dialog box opens. The **Search From** field is populated with the **Committed Transactions** folder.
4. Enter search criteria and click **Find**.

To search committed transactions from the Find dialog box:

1. From the **Tools** menu, select **Find**.
2. Click **Browse**.
3. Navigate to the **Committed Transactions** folder and select it.
4. Enter search criteria and click **Find**.

Exporting Transactions

To export all transactions:

1. From the **Tools** menu, select **Options**.

   The Transactions dialog box opens.
2. Select the **Transactions** property page.
3. Click **Export**.

   The transactions are exported to a text file. The Save As dialog box opens.
4. Enter the file name for the destination text file and click **Save**.

To export a single transaction:

1. Locate the single transaction that you want to export.
2. Right-click the transaction and select **Properties**.

   The Transaction dialog box opens.
3. Select the **Transactions** property page.
4. Click **Export**.

   The transaction is exported to a text file. The Save As dialog box opens.
5. Enter the file name for the destination text file and click **Save**.

Viewing a List of Transactions

IP Service Activator lists the transactions that are currently held in the transaction store on the **System** tab under the **Transactions** folder.

Depending on your organization, the transactions that are listed may include:

- Transactions created by you
- Transactions created by another user working remotely

- Transactions created by a third-party tool through the OSS Integration Manager (OIM)

> **Note:** The number of transactions that are listed in the Committed Transactions folder is configurable. See *IP Service Activator System Administrator's Guide* for details.

To view a list of transactions, on the **System** tab, open the **Transactions** folder.

Pending transactions are listed under the **Transactions** folder, and scheduled and committed transactions are grouped into subfolders.

Viewing Transaction Details

You can view detailed information about all transactions in the Details pane and sort by any of the displayed columns. If you have adopted a naming convention for transactions based on the first letters of the transaction name, this is a way of sorting related transactions. See "Changing Views" for information about sorting by a column.

To view transaction details:

1. On the **System** tab, double-click the **Transactions** folder.

   The Details pane lists transaction information under the following headings:

   - **Name**: Identifier for the transaction
   - **Description**: Additional transaction information (optional)
   - **State**: Transaction status (see "Transaction State")
   - **Schedule**: Time at which the transaction was or will be committed
   - **Username**: Name of the user who created the transaction, entered automatically by IP Service Activator.
   - **Size**: Size of the transaction in bytes
   - **System Configuration**: Number of concrete objects created by the transaction and their status.
   - **Fault Level**: Number and type of faults created by the transaction.
   - **Owner**: If ownership of the transaction has been specified, the owner's user name.
   - **OwnerGroup**: If ownership of the transaction has been specified, the group to which the owner belongs.

Transaction State

Every transaction has a state that indicates where it is in the transaction lifecycle. Table 8–1 describes these states.

Table 8–1 Transaction States

State	Description
Pending	A transaction that has been saved to the transaction store but whose changes have not been propagated to the network. Pending transactions are visible in all user interfaces.
Scheduled	A transaction that will be committed at a future date and time.

Table 8–1 (Cont.) Transaction States

State	Description
Failed	A scheduled transaction that IP Service Activator was unable to commit due to conflicts with the common object model.
Committed	A transaction whose changes have been saved to the database and related configuration changes propagated to the network.

Checking the Origin of Transactions

To check which user created a transaction, you can:

- Open the relevant transaction's properties dialog box and check the **Username** field on the **Transaction** property page.
- List all transaction details in the Details pane (see "Viewing Transaction Details").
- Check the audit trail (viewable under the **System** tab's **System Logs** folder).

Selecting Transactions

You can select a single transaction or multiple transactions listed in the Details pane and select an action to perform on them. For example, you can merge a number of transactions in one step by multi-selecting them and selecting **Merge** from the context menu. Where several transactions are selected, IP Service Activator processes them in the order in which they are selected.

See "Selecting Objects" for information about selecting objects in the Details pane.

Provisioning Status

All changes in IP Service Activator are transactions. The transaction entry, as a whole, has a provisioning status that you can view. Provisioning status applies to a transaction only if transaction status monitoring is enabled. See "Enabling Transaction Status Monitoring" for information about enabling transaction status monitoring.

To view provisioning status:

1. From the **System** menu, select **Committed Transactions**, then select **Transaction Entry**, and then select **Properties**.

   The Transaction dialog box opens.

Table 8–2 lists possible **Provisioning Status** values and their meaning.

Table 8–2 Provisioning Status Values

Provisioning Status Value	Meaning
Pending	No confirmation from the downstream components that the transaction was received
Succeeded	Confirmation received indicating successful creation, modification, or uninstallation of corresponding configuration
Failed	Confirmation received indicating the creation, modification, or uninstallation of the corresponding configuration change was rejected
Timedout	Confirmation not received during the timeout period specified by **ConcreteActivationStatusTimeout** (default: 30 minutes)

Status of Concretes within a Transaction

You can view the concrete activation status in the client. Valid activation status values and their meanings are in Table 8–3.

Table 8–3 Activation Status Values

Activation Status Value	Meaning
Inactive	A concrete is inactive as long as the transaction is not yet committed to the policy server
Pending	A concrete is pending when notification has not yet been sent
Succeeded	Notification is received that a concrete operation succeeded
Failed	Notification is received that a concrete operation failed
Timedout	Notification is received that a concrete operation timed out
TestSucceeded	Notification is received that a concrete operation succeeded, when the operation is run in offline, test mode
TestFailed	Notification is received that a concrete operation failed, when the operation is run in offline, test mode
Conflicted	Notification is received of a conflict that must be resolved

Enabling Transaction Status Monitoring

Transaction status monitoring allows IP Service Activator to provide feedback regarding the success or failure of operations that result in deleted concretes.

Use the command line to configure the following component parameters:

- **Full transaction monitoring**: controls whether the policy server sets the provisioning status and reason for failure of completed transactions
- **Timeout interval**: represents the maximum time a transaction is allowed to stay in Pending state awaiting completion notifications for its concretes
- **Fault wait interval**: represents the time that the policy server waits before collecting the faults for a failed transaction
- **Transaction monitoring poll interval**: represents the frequency at which the policy server checks each in-flight transaction to determine if transactions have timed out

> **Note:** These parameters apply only to the policy server.

Table 8–4 shows the transaction monitoring command line and component parameters.

Table 8–4 Transaction Monitoring

Command Line Parameter	Component Parameter	Values	Range (Default)
FullTransactionMonitoring	TransactionMonitoring.FullMonitoring	enabled/disabled	Enabled
TransactionMonitoringTimeout	TransactionMonitoring.Timeout	seconds	60-14400 (900)
FaultWaitTime	TransactionMonitoring.FaultWaitTime	seconds	10-600 (10)

Table 8–4 (Cont.) Transaction Monitoring

Command Line Parameter	Component Parameter	Values	Range (Default)
TransactionMonitoring PollInterval	TransactionMonitoring.PollInterval	seconds	10-3600 (60)