20 Managing AutoTracker VLANs

In a large, flat, switched network, broadcast traffic can overload a network based primarily on port-based Groups. Through the use of AutoTracker VLANs, you can control broadcast traffic such that it is forwarded only to those VLANs where it needs to be sent.

VLANs are created within a Group to subdivide network traffic based on specific criteria. The criteria you use to define a VLAN are called AutoTracker policies. AutoTracker policies can be defined by port, MAC address, protocol, network address, a user-defined policy, or a multicast policy. You can also define multiple policies—also referred to as "rules"—for a VLAN if you wish. A port or device is included in a VLAN if it matches any one VLAN rule. For example, you can define rules based on MAC address and rules based on protocol in the same VLAN.

A Group defines a physical space within the network—a set of ports. The policies that you define for VLAN membership are applied to all traffic on those ports, but not to traffic on ports outside the Group.

You can create two types of policy-based VLANs: AutoTracker VLANs and multicast VLANs. You can create up to 31 AutoTracker VLANs and up to 32 multicast VLANs in any one Group. AutoTracker VLANs and multicast VLANs operate independently of one another: the policies you establish for AutoTracker VLANs neither conflict nor interfere with the policies you establish for multicast VLANs, even when those policies involve the same ports or MAC addresses.

This chapter provides an overview of AutoTracker VLANs and multicast VLANs as well as instructions for managing and monitoring each type of VLAN. Instructions for configuring AutoTracker policies can be found in Chapter 18, "Configuring Group and VLAN Policies."

The AutoTracker Menu

All software commands for configuring AutoTracker policies and AutoTracker/multicast VLANs are in the AutoTracker menu. This menu is a submenu of the VLAN menu. You can access the AutoTracker menu by typing **at** at any prompt. The menu displays as follows:

```
Command      Auto-Tracker Management Menu
-----------  ------------------------------------------------------------
cratvl       Create an Auto-Tracker VLAN
atvl         View definition of Auto-Tracker VLAN
viatrl       View Auto-Tracker Rule Configuration
rmatvl       Delete an Auto-Tracker VLAN
modatvl      Modify definition of an Auto-Tracker VLAN
vivl         View list of Active Auto-Tracker VLANs on an interface
fwtvl        View VLAN assignment of learned MAC addresses
defvl        Enable or disable membership in default VLAN
crmcvl       Create a Multicast VLAN
mcvl         View definition of Multicast VLAN
vimcrl       View Multicast VLAN Rule Configuration
rmmcvl       Delete a Multicast VLAN
modmcvl      Modify definition of a Multicast VLAN
vimcvl       View list of Active Multicast VLANs on an interface
gmstat       Turn Group Mobility Status ON or OFF
vpl          View Virtual Ports in a Mobile Group
vigl         View Mobile Group List for a Virtual Port
cats         Create Auto-Activated Services
data         Delete Auto-Activated Services
vats         View Auto-Activated Services
vag          View Authenticated Groups
gmcfg        Configure Group Mobility Parameters
mag          Modify Authenticated Group
xip          Enter the Xylan Inter-switch Protocol (XIP) sub-menu

     Main        File       Summary     VLAN       Networking
     Interface   Security   System      Services   Help
```

The commands on the AutoTracker menu can be roughly divided into two halves. The first half of commands—listed from **cratvl** to **vimcvl**—apply mainly to AutoTracker VLANs (i.e., VLANs created inside non-mobile groups). An exception to this rule is the **modatvl** command, which can be used to modify AutoTracker policies for VLANs or mobile groups. In addition many of the informational commands apply to both VLANs and mobile groups. The commands that apply to AutoTracker VLANs are described in this chapter. Multicast VLANs are described in Chapter 21, "Multicast VLANs." The **mag** command is described in the *Switched Network Services User Manual*. The XIP sub-menu is described in Chapter 19, "Inter-switch Protocols."

The commands from **gmstat** to **gmcfg** apply strictly to mobile groups. All of the commands in this second set are described in Chapter 17, "Managing Groups and Ports."

AutoTracker VLANs

AutoTracker VLANs enable you to control communications between end stations in your network. You define policies that determine membership in the VLAN and AutoTracker automatically locates ports or devices within the Group that fit the policies and places them into the VLAN.

You can define physical policies or logical policies (or combinations thereof) to determine membership in AutoTracker VLANs. Physical policies consist of port rules: you define the VLAN members as one or more specific ports and VLAN membership is limited to the ports defined and the MAC addresses of devices connected to those ports.

Logical VLAN policies can consist of MAC address rules, protocol rules, network address rules, or user-defined rules. Ports are assigned to VLANs that have logical rules when the MPM examines frames that originate from devices connected to the Group's set of ports. If a frame is received that matches a logical VLAN rule, the source device's MAC address and the port to which the source device is connected are both made VLAN members.

The members of an AutoTracker VLAN thus consist of source devices originating frames that fit the VLAN's policies and the ports to which those source devices are connected. Instructions for creating AutoTracker VLANs begin on page 20-3.

AutoTracker VLAN Policies

You can define a maximum of 32 AutoTracker policies of each type per Group. There is no restriction on the number of rules you can define per VLAN, as long as the maximum number of policies for the Group is not exceeded.

A switch port – or a device connected to a switch port – can belong to more than one VLAN simultaneously, as determined by the rules the port or device matches. A port or device is included in a VLAN if it matches any one rule.

You can define the following types of rules:

Port Policies. Port policies enable you to define membership in the VLAN on the basis of ports. Members of the VLAN will consist of devices connected to specific ports on one switch or on multiple switches in the Group.

MAC Address Policies. MAC address policies enable you to define membership in the VLAN on the basis of devices' MAC addresses. This is the simplest type of rule and provides the maximum degree of control and security. Members of the VLAN will consist of devices with specific MAC addresses. These devices may all be connected to one switch or they may be connected to different switches in the Group. A maximum of 10,240 MAC addresses are supported per policy.

Protocol Policies. Protocol policies enable you to define membership in the VLAN on the basis of the protocol that devices use to communicate. All devices that communicate with the specified protocol become members of the VLAN.

You can specify VLAN membership according to the following protocols: IP, IPX, AppleTalk, or DECNet. In addition, you can specify membership according to Ethernet type, source and destination SAP (service access protocol) header values, or SNAP (sub-network access protocol) type.

Network Address Policies. Network address policies enable you to define membership in the VLAN on the basis of network address criteria.

For example, you can specify that all IP users with a specific subnet mask be included in the VLAN. Or, you can specify that all IPX users in a specific network address area using a certain encapsulation type be included in the VLAN.

If you define network address and port or protocol rules in the same VLAN, the network address rules will take precedence over the port and protocol rules should any conflict arise. To reverse this precedence (i.e., port and protocol rules take precedence over network address rules) you must add the following line to the switch's **mp4.cmd** file:

```
Precedence=0
```

User-Defined Policies. User-defined policies enable you to define membership in the VLAN on the basis of a specific pattern within a frame. All devices that originate frames containing this pattern are assigned to the VLAN. The pattern is specified by defining an offset, a value, and a mask.

Port Binding Policies. A port binding policy specifies a particular device to be included in the mobile group or AutoTracker VLAN. You can bind a device's IP address to a switch port and a MAC address, or bind a device's MAC address to a protocol and a switch port.

DHCP Port Policies. These policies are similar to standard port policies, but apply to switch ports to which DHCP client workstations are attached.

DHCP MAC Address Policies. These policies are similar to standard MAC address policies, but apply to the MAC addresses of DHCP client workstations only.

The Default VLAN

The default AutoTracker VLAN, also referred to as VLAN #1, is different from other AutoTracker VLANs. The following list outlines some of these differences.

1. The default VLAN is automatically created when you create a new Group. Non-default VLANs must be created through the **cratvl** command.

2. The default VLAN cannot be removed. Other VLANs can be removed through the **rmatvl** command.

3. You cannot apply AutoTracker policies to the default VLAN. Other non-default AutoTracker VLANs allow you to apply any policy to them.

You can enable routing on the default VLAN. You enable the default VLAN virtual router through the **crgp** or **modvl** command. See Chapter 17, "Managing Groups and Ports," for further information on the virtual router port on the default VLAN.

All ports and devices in a Group initially belong to default VLAN #1. All physical switch ports always remain members of the default VLAN, but they can also become members of other VLANs. It is not possible to delete a physical switch port from VLAN #1. Individual network devices, however, can move out of VLAN #1. All MAC devices are also initially part of default VLAN #1. However, when a MAC device is removed from default VLAN #1 and moved into a non-default VLAN, it is deleted from default VLAN #1.

The default VLAN is explained further in other sections of this chapter. See *How Devices are Assigned to AutoTracker VLANs* on page 20-5 for a discussion of default VLAN membership issues and the **defvl** command. Also, see *Application Example 4* in Chapter 22, "AutoTracker VLAN Application Examples," for discussions of routing issues and the default VLAN.

How Devices are Assigned to AutoTracker VLANs

When a broadcast frame, a multicast frame, or a unicast frame from an unknown device is received at a switching module, the frame is forwarded to the MPM for processing. Source learning logic on the MPM module examines the entire frame to determine the VLAN or VLANs in which the originating device should be a member. If the frame matches any one policy defined for a VLAN, the originating device (and the port to which it is connected) are made members of that VLAN. If the frame does *not* match any VLAN policy, one of the following occurs:

- If the **defvl** command is on, the source device is made a member of Default VLAN #1 in the Group of which the source port is a member. The **defvl** command determines whether traffic from devices that do not match any policies is assigned to the default VLAN or dropped. (See "The defvl Command" below for more information on this command.)
- If the **defvl** command is off, all traffic from the source device is dropped.

Please Take Note

A broadcast or multicast frame is processed to determine the source device's VLAN membership each time it is received. A unicast frame is processed to determine the source device's VLAN membership only the first time it is received.

When the MPM module has determined the VLAN or VLANs in which the originating device belongs, it relays this information to the switching module. The switching module updates a VLAN membership flag attached to the frame's source MAC address in the CAM (content-addressable memory). The frame is then switched based on this membership flag.

Refer to Chapter 22, "AutoTracker VLAN Application Examples," for information on AutoTracker VLAN assignments in specific network situations.

The defvl Command

You can turn the **defvl** command on and off simply by entering **defvl on** or **defvl off**. If you enter the command without any parameters, it displays the current setting for the Default VLAN. For example, if source devices are automatically placed in the Default VLAN when they do not match any VLAN policy rule, the following message would display:

```
membership in default vlan is currently on
```

If source devices are automatically dropped when they do not match any VLAN policy, the following message would display:

```
membership in default vlan is currently off
```

The defvl command applies to all Groups in an OmniStack and it is only applicable if there is at least one AutoTracker rule configured.

Devices that Generate a Secondary Traffic Type

Source devices sometimes generate more than one traffic type; for example, a device could generate IP traffic primarily but also generate a secondary stream of AppleTalk. When a device generates secondary traffic that does not match any existing VLAN policy, that traffic is grouped into the primary VLAN of which the device is a member.

To continue the example, if a device generates both IP and AppleTalk, and both an IP VLAN and an AppleTalk VLAN exist, that device is made a member of both VLANs and no problem occurs. If, however, an AppleTalk VLAN does not exist, all traffic from that device is grouped into the existing VLAN of which the device is a member – in this example, the IP VLAN. This can cause communication problems, as explained below. ***For this reason, it is advisable to create VLANs that accommodate all known network traffic.***

In this example Device A is assigned to default VLAN #1 because it does not match any existing VLAN policy.

Devices B and C are assigned to VLAN 2 because they generate IP traffic. The secondary AppleTalk traffic Device C generates is also grouped into VLAN 2, since the AppleTalk traffic does not match any existing VLAN policy.

The result is that Devices A and C are unable to communicate.

Creation of an AppleTalk protocol VLAN solves this problem. If an AppleTalk VLAN exists, Device A will be assigned to it and removed from Default VLAN #1. Device C will be assigned to both the IP VLAN and the AppleTalk VLAN. Devices A and C can then communicate.

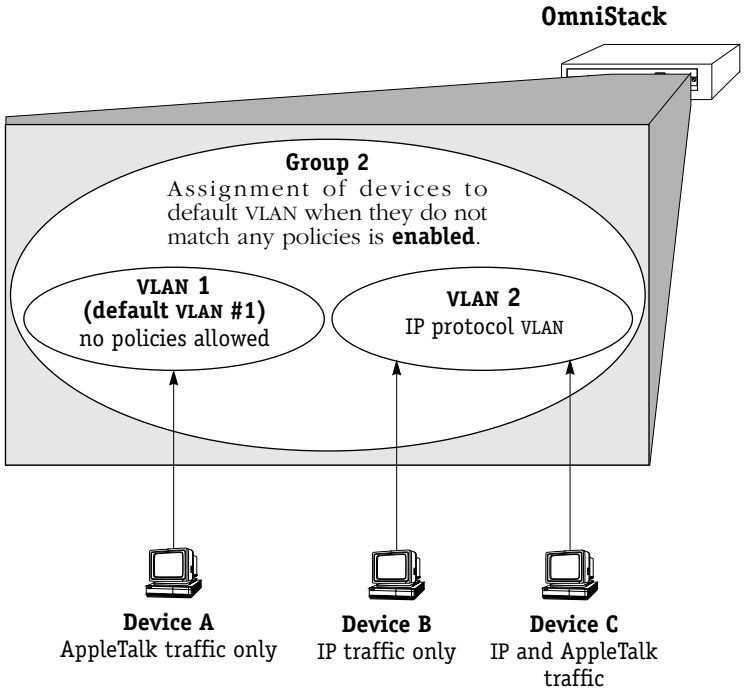

How Devices are Assigned to AutoTracker VLANs *(continued)*

Router Traffic in IP and IPX Network Address VLANs

Prior to release 2.1, AutoTracker handled VLAN assignments for router traffic in IP and IPX network address VLANs in the same manner as normal traffic. In release 2.1 and later, AutoTracker differentiates router traffic from normal traffic and can distinguish traffic that is routed *through* a router from traffic that is generated *by* a router.

AutoTracker now determines VLAN assignments for router interfaces (that is, the MAC addresses of router interface ports) in IP and IPX network address VLANs based on router update messages generated by the router itself. This minimizes VLAN leakage and avoids the problem situation described on the facing page.

The Problem with Router Traffic

AutoTracker functions on the assumption that data in a frame can be associated with the frame's source MAC address. For example, if a frame has an IPX network number of 300, AutoTracker assumes that it has received the frame directly and that the source device is a member of IPX network 300. This is not true in the case of routed frames. Routers route frames from one network to another by changing the frame's MAC header but keeping the layer 3 content intact. This can lead to the problem situation described on the facing page.

In the network on the facing page, Device A gets correctly assigned to VLAN 2 and Device B gets correctly assigned to VLAN 3 without problem. The two router interfaces will be assigned to the correct VLANs *if AutoTracker learns the router interface MAC addresses from their RIP updates*. However, this may not happen. The problem situation on the facing page shows what can occur if AutoTracker learns the router interface MAC addresses from traffic routed through the router rather than from traffic generated by the router (such as a RIP update).

How AutoTracker Handles Router Traffic

To avoid the problem situation on the facing page, AutoTracker now determines if any IP or IPX device it has learned is a router. If it is, AutoTracker marks the device as a router, unlearns all previous VLAN assignments for that device, and reassigns the device based on a router-generated update packet (such as a RIP packet).

AutoTracker determines if a learned device is a router by searching further within the frame. For example, if AutoTracker receives an IP frame, it searches beyond the source IP address and also checks if the IP frame is a RIP, OSPF, BGP, DVMRP, or IGRP update. If it is, as explained, AutoTracker marks the device as a router, unlearns its previous VLAN assignments, and reassigns it using the router-generated update packet.

AutoTracker recognizes the following types of router-generated frames:

- IP protocol: RIP frames, OSPF frames, BGP4 frames, DVRMP frames, and IGRP frames
- IPX protocol: IPX RIP frames and SAP frames

AutoTracker maintains a record of the devices it has learned are routers. Each time a router-generated frame is received from a device marked as a router, AutoTracker updates that device's membership in IP or IPX network address VLANs. If a frame received from a device marked as a router is not IP or IPX, VLAN membership is updated normally.

Please Take Note

This special handling of router traffic occurs in IP and IPX network address VLANs only. Note that it does not alter normal VLAN assignment processes such as checking for VLAN policy matches other than IP or IPX network address.

How Routed Frames can Confuse VLAN Assignment

❷ The router receives the frame on the interface for Network 2 and routes the frame to the interface for Network 3. To do this, the router strips the MAC header from the frame and inserts the MAC address of its interface for Network 3. The frame now specifies its source as Network 2, MAC address Y.

Frame
from Network 2, MAC address Y

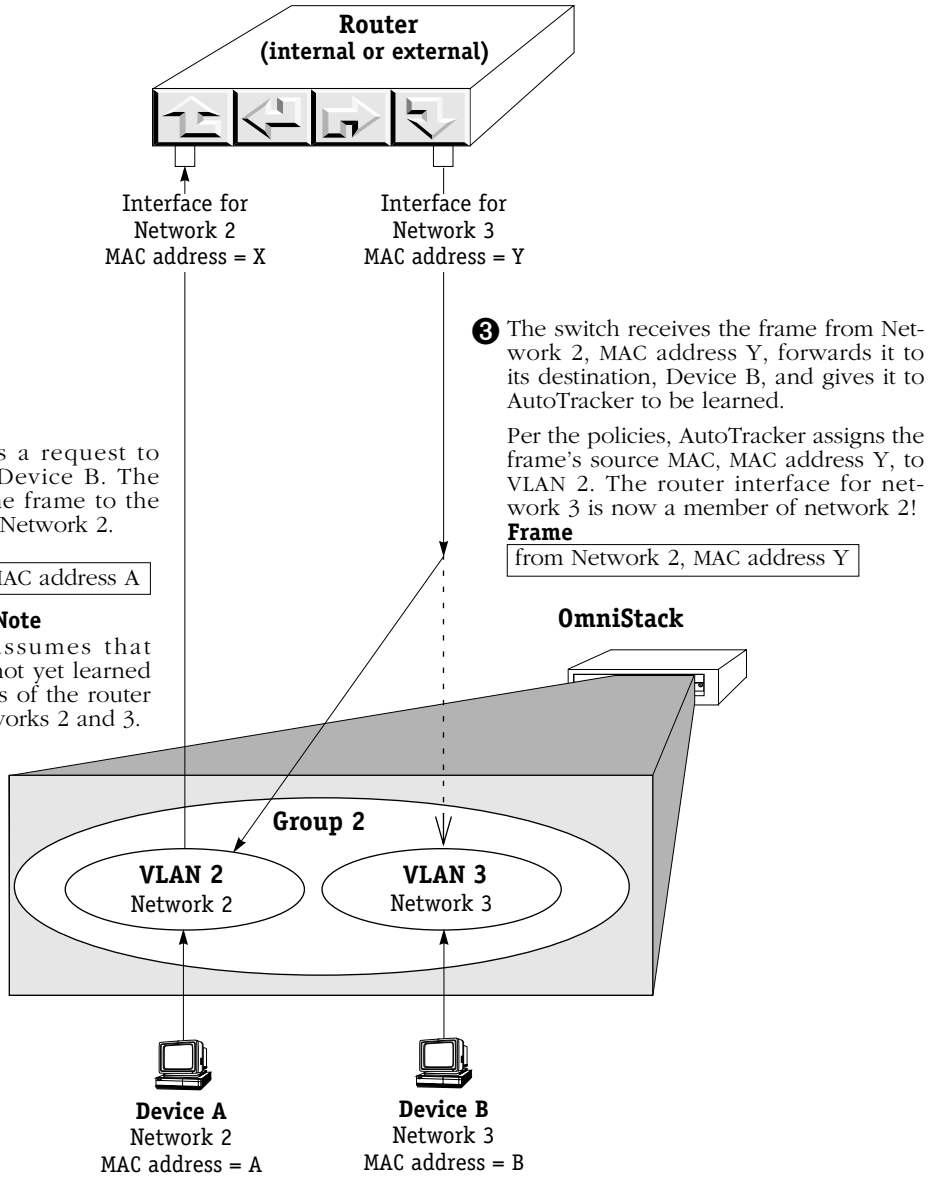

❸ The switch receives the frame from Network 2, MAC address Y, forwards it to its destination, Device B, and gives it to AutoTracker to be learned.

Per the policies, AutoTracker assigns the frame's source MAC, MAC address Y, to VLAN 2. The router interface for network 3 is now a member of network 2!

Frame
from Network 2, MAC address Y

❶ Device A initiates a request to route a frame to Device B. The switch forwards the frame to the router interface for Network 2.

Frame
from Network 2, MAC address A

Please Note
This example assumes that AutoTracker has not yet learned the MAC addresses of the router interfaces for Networks 2 and 3.

❹ Let's say that the next transmission is a RIP update from the router interface for network 3. The source of the RIP update is Network 3, MAC address Y. AutoTracker thus assigns MAC address Y to VLAN 3. MAC address Y is now assigned to both VLAN 2 and VLAN 3.

The same situation can occur with MAC address X on the router interface for network 2. Both router interfaces will be members of both VLANs and will transmit RIP updates to both.

If this is an IPX network and IPX servers are members of these VLANs, they will respond with router configuration errors. If this is an IP network and devices A and B are IP workstations listening to RIP, they will respond with invalid network address errors.

How Devices are Assigned to AutoTracker VLANs *(continued)*

Port Policy Functionality

In release 2.1 and later, AutoTracker's VLAN port policy can be set to operate in either of two distinct modes:

- In the original mode, wherein membership in all VLANs active on a port is inherited by all devices connected to that port. Original port policy functionality is explained on page 20-10.
- In a new mode, wherein membership in all VLANs active on a port **is not** inherited by all devices connected to that port. This is the current, default functionality with which the switch ships. Current port policy functionality is explained on page 20-11.

Port policy functionality is set on a switch-wide basis, via a flag in the switch's **mp4.cmd** file called **reg_port_rule**. The OmniStack ships with port policy functionality set to operate in the new mode. You can revert the switch to original port policy functionality by editing the **mp4.cmd** file and setting the **reg_port_rule** flag to 1. You must then restart the switch. (The **mp4.cmd** file is accessed, and can be edited, via the switch User Interface. You can view the current setting of **reg_port_rule** with the **view mpm.cmd** command. See Chapter 5, "Managing Files," for information on editing the **mp4.cmd** file.)

Why the New Functionality?

Port policies can cause problems in a multi-switch environment. AutoTracker assumes that each switch in a multi-switch environment can independently arrive at identical VLAN assignments for all devices in the network. This is not true when port policies are in effect because of their very nature: port policies are switch-specific and not network wide. The figure on page 20-10, which explains original port policy functionality, provides an example of how port policies can result in inconsistent VLAN membership between two switches – notice the inconsistent VLAN membership in OmniStack 1 and in OmniStack 2.

The use of port policies in a multi-switch environment can result in connectivity problems if the source switch and the destination switch are separated by other switches. The switches along the path of the frame will not have identical VLAN memberships. At any particular switch along the path, frames could be lost because of inconsistencies in the VLAN membership of the frames' source and destination devices.

In addition, AutoTracker maintains devices in the same VLAN without regard to the devices' location – provided the devices match the same AutoTracker policies throughout the network. Multiple switches will assign a device to the same VLANs provided that device matches the same policies on each switch. This is not possible when port policies are in effect because, as stated, by their very nature port policies are switch-specific and not network-wide.

For these reasons, the OmniStack now ships with new port policy functionality (although, as explained, you can revert the switch to original port policy functionality if you wish). The new functionality still enables users to assign ports to VLANs and still enables those ports to carry traffic for those VLANs. However, with the new functionality, port policies are not used to learn VLAN assignments for traffic received on ports (as explained on page 20-11). In order for a device to be assigned to a VLAN, it must match an existing logical policy of the VLAN. This is explained on page 20-13.

The Following Examples

The following pages provide examples of original and current port policy functionality. The limitations of port policies become apparent if one tries to use port policies to create two VLANs in these sample networks, one for Devices A and B and one for Devices C and D.

Original Port Policy Functionality (reg_port_rule = 1)

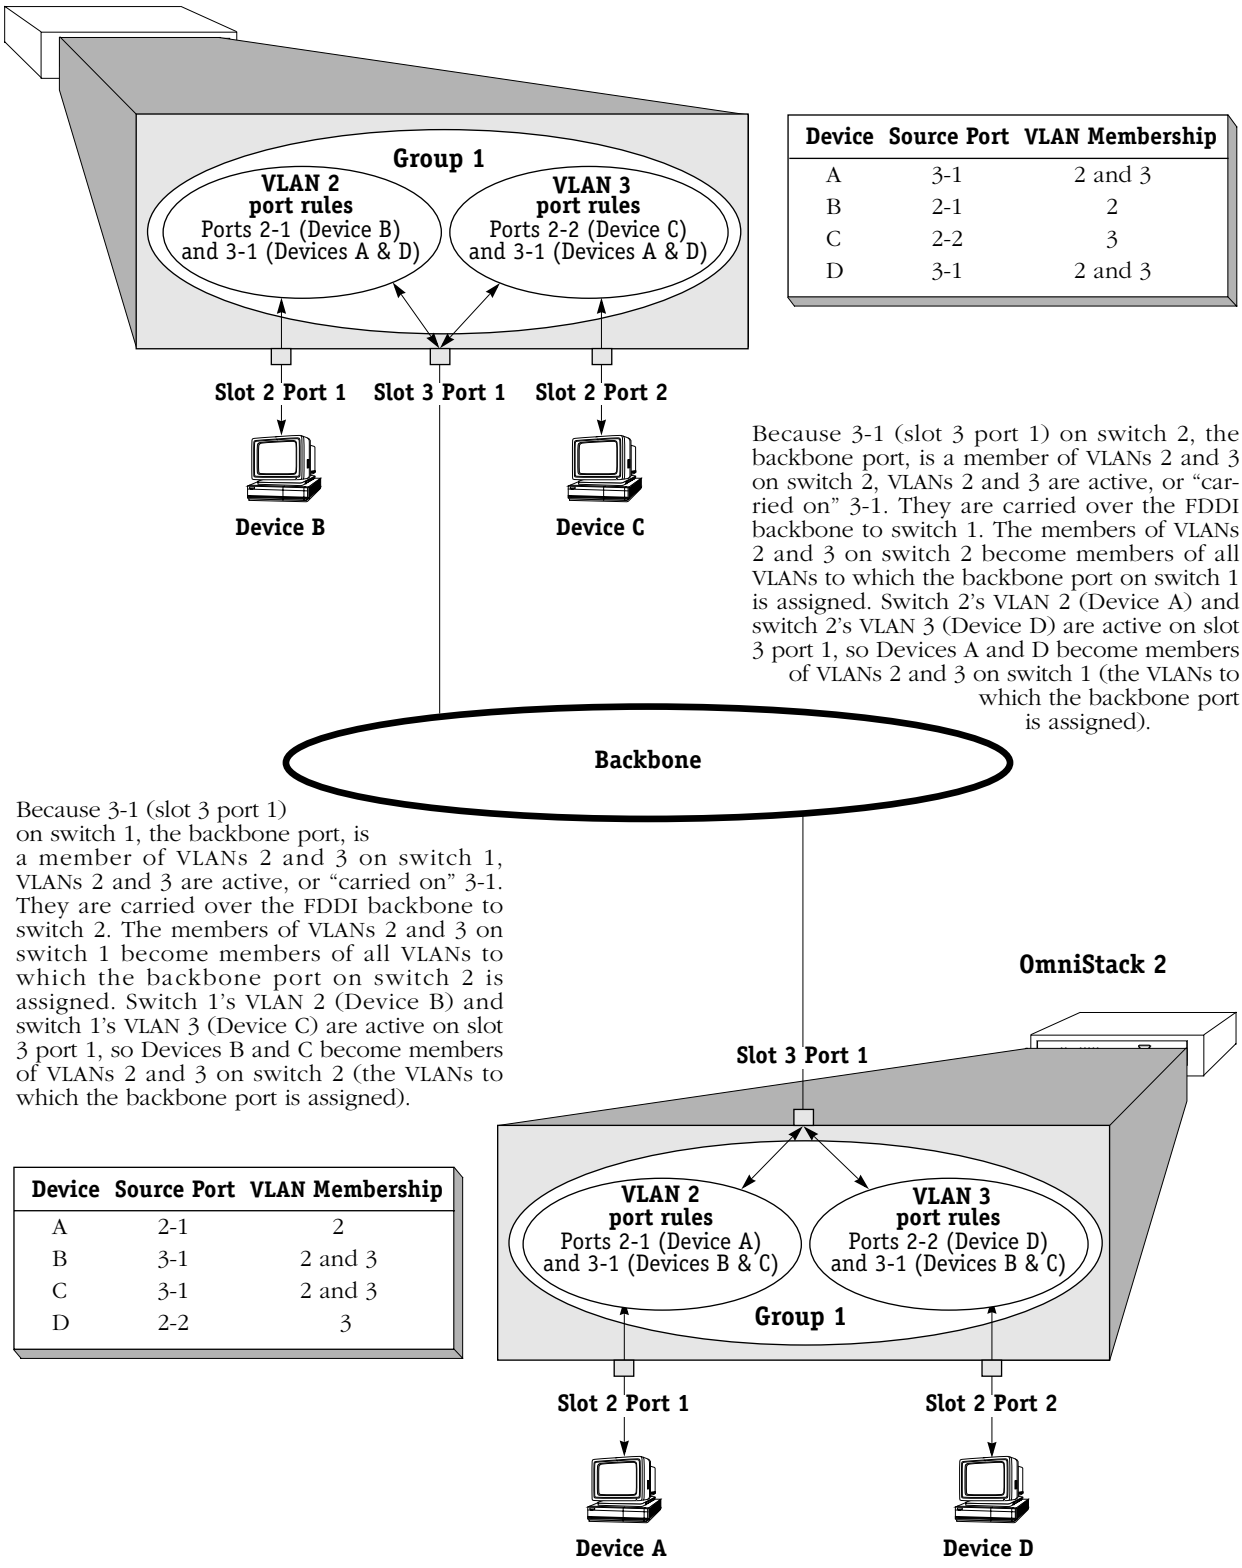

Device	Source Port	VLAN Membership
A	3-1	2 and 3
B	2-1	2
C	2-2	3
D	3-1	2 and 3

Because 3-1 (slot 3 port 1) on switch 2, the backbone port, is a member of VLANs 2 and 3 on switch 2, VLANs 2 and 3 are active, or "carried on" 3-1. They are carried over the FDDI backbone to switch 1. The members of VLANs 2 and 3 on switch 2 become members of all VLANs to which the backbone port on switch 1 is assigned. Switch 2's VLAN 2 (Device A) and switch 2's VLAN 3 (Device D) are active on slot 3 port 1, so Devices A and D become members of VLANs 2 and 3 on switch 1 (the VLANs to which the backbone port is assigned).

Because 3-1 (slot 3 port 1) on switch 1, the backbone port, is a member of VLANs 2 and 3 on switch 1, VLANs 2 and 3 are active, or "carried on" 3-1. They are carried over the FDDI backbone to switch 2. The members of VLANs 2 and 3 on switch 1 become members of all VLANs to which the backbone port on switch 2 is assigned. Switch 1's VLAN 2 (Device B) and switch 1's VLAN 3 (Device C) are active on slot 3 port 1, so Devices B and C become members of VLANs 2 and 3 on switch 2 (the VLANs to which the backbone port is assigned).

Device	Source Port	VLAN Membership
A	2-1	2
B	3-1	2 and 3
C	3-1	2 and 3
D	2-2	3

Current Port Policy Functionality (reg_port_rule = 0)

OmniStack 1

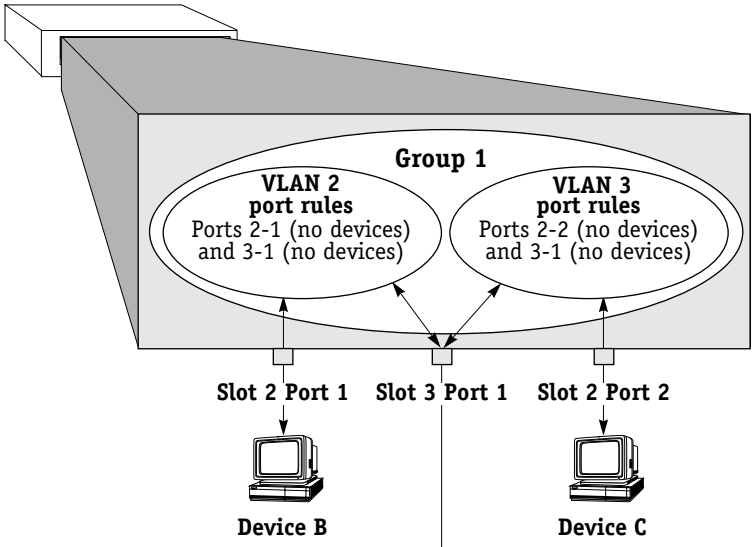

Device	Source Port	VLAN Membership
A	3-1	default VLAN #1
B	2-1	default VLAN #1
C	2-2	default VLAN #1
D	3-1	default VLAN #1

With current port policy functionality, VLANs are not active, or "carried on" ports. Port 3-1 (slot 3 port 1) on switch 2, the backbone port, is still a member of VLANs 2 and 3 on switch 2, but VLANs 2 and 3 **are not** carried over the FDDI backbone to switch 1. The members of VLANs 2 and 3 on switch 2 **do not** become members of all VLANs to which the backbone port on switch 1 is assigned. Rather, they become members of default VLAN #1 on switch 1 if they do not match any of switch 1's existing VLAN policies and if **defvl** is on for the switch. If **defvl** is off, the traffic is dropped.

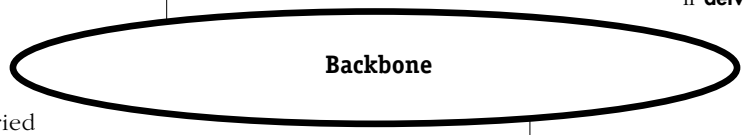

With current port policy functionality, VLANs are not active, or "carried on" ports. Port 3-1 (slot 3 port 1) on switch 1, the backbone port, is still a member of VLANs 2 and 3 on switch 1, but VLANs 2 and 3 **are not** carried over the FDDI backbone to switch 2. The members of VLANs 2 and 3 on switch 1 **do not** become members of all VLANs to which the backbone port on switch 2 is assigned. Rather, they become members of default VLAN #1 on switch 2 if they do not match any of switch 2's existing VLAN policies and if **defvl** is on for the switch. If **defvl** is off, the traffic is dropped.

Device	Source Port	VLAN Membership
A	2-1	default VLAN #1
B	3-1	default VLAN #1
C	3-1	default VLAN #1
D	2-2	default VLAN #1

OmniStack 2

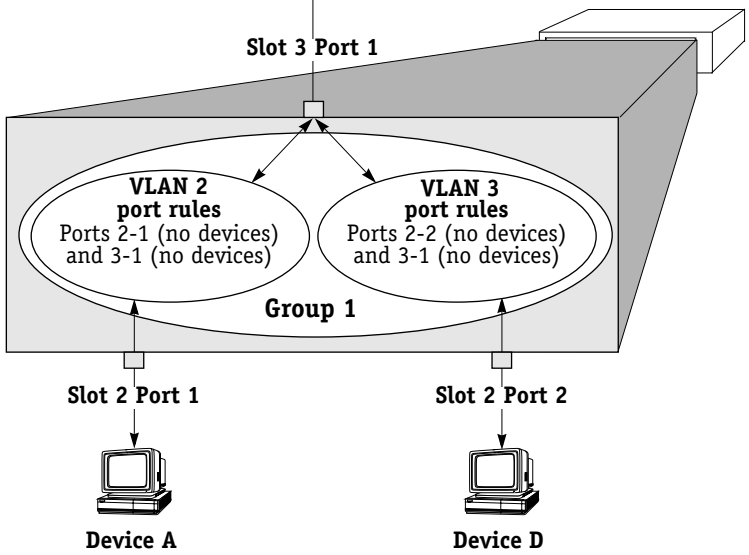

The Usefulness of Port Policies

As has been explained – and as illustrated on page 20-10 – original port policy functionality is not well-suited to the creation of consistent VLAN membership in a multi-switch environment. Current port policy functionality – as illustrated on page 20-11 – neither contributes to nor participates in VLAN assignments. Port policies, either original or current, are in fact not useful in the creation of consistent VLAN membership across multiple switches. Logical policies are of far greater use, as illustrated on page 20-13. So, why use port policies at all?

Port Policies are Useful in these Situations:

- **Silent stations**. If a device does not transmit traffic (such as a printer), the port to which the device is connected never gets assigned to VLANs. It is then impossible for other stations to communicate with that device. Creating a port policy that assigns the silent device's port to one or more VLANs will enable traffic to flow out that port to the silent device.
- **Inactive VLANs**. AutoTracker does not activate a VLAN – or its internal router – until a port is assigned to that VLAN. AutoTracker assigns ports to VLANs with port policies immediately. However, AutoTracker only assigns ports to VLANs with logical policies when a frame is received from a source device that matches the VLAN's policies. This means that, in some network situations, you may need to assign a port policy to a VLAN to force it active. *Application Example 4* in Chapter 22, "AutoTracker VLAN Application Examples," provides an example of this.
- **Backbone connections**. A port policy that assigns the backbone port to a VLAN will enable traffic from that VLAN to flow out onto the backbone.

♦ Important Note ♦

If you are using port policies to extend VLANs across a backbone, you are strongly advised to use current (default) port policy functionality. If you use original port policy functionality, you are, in effect, placing all devices learned from the backbone port into the same VLAN. If the port policy is configured for all VLANs (so that all VLANs can communicate over the backbone), all devices learned from the backbone port are assigned to all VLANs. This is not desirable – it would subject locally-connected devices to all the backbone traffic.

So How Do I Get Devices Assigned to VLANs Over a Backbone?

The way to get devices assigned to VLANs over a backbone is to define logical VLAN policies that so assign them. An example is shown on the facing page utilizing IP and IPX protocol policies. The network on the facing page uses port policies (and current port policy functionality) to assign the backbone port to VLANs on each switch so that traffic can flow out onto the backbone from these VLANs.

The problem of remote VLAN assignments is solved by the IP and IPX protocol policies. When a frame is received from a backbone port, the frame is examined to determine if it matches any VLAN membership rules. Let's say Device D on switch 2 transmits an IP frame. The frame travels the FDDI backbone and enters switch 1 on port 3-1. AutoTracker learns the frame and assigns it to VLAN 3, since VLAN 3 has an IP protocol policy and the frame is IP.

Notice that with this approach:

- VLAN membership is consistent between the two switches.
- In a multi-switch environment, no frames are lost in switches along the traffic path because of the inconsistent VLAN membership of a frame's source and destination devices.
- Devices can be moved from switch to switch and they will be assigned to the same VLAN – without reconfiguring AutoTracker or the device.
- As was the original intent, it is possible to create two VLANs in this sample network, one for Devices A and B and one for Devices C and D. As is apparent, this was impossible using port policies.

An Example of VLAN Assignment Using Logical Policies and Current Port Policy Functionality (reg_port_rule = 0)

OmniStack 1

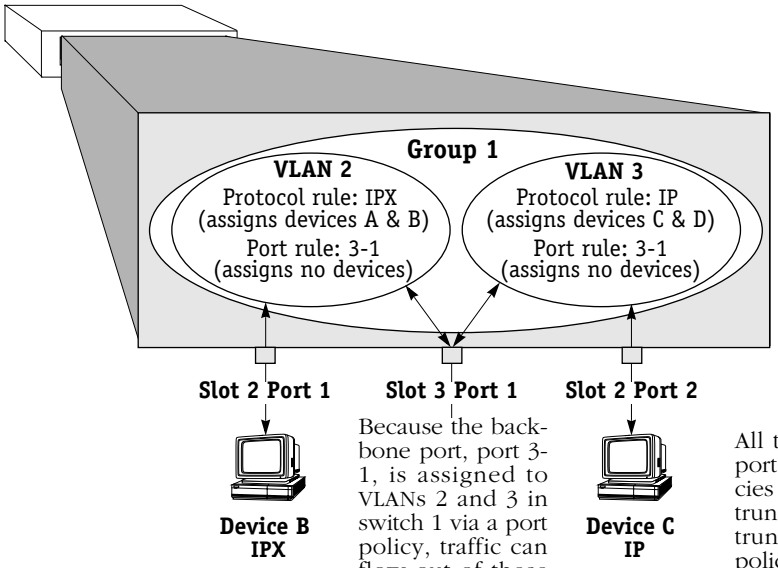

Device	Source Port	VLAN Membership
A	3-1	2
B	2-1	2
C	2-2	3
D	3-1	3

Because the backbone port, port 3-1, is assigned to VLANs 2 and 3 in switch 1 via a port policy, traffic can flow out of these VLANs onto the backbone.

All traffic entering switch 1 from the backbone port, port 3-1, is checked against switch 1's policies for VLAN membership in the Group(s) being trunked. In this example only Group 1 is being trunked, so traffic is checked against the VLAN policies defined for Group 1's VLANs. If you want devices from switch 2 included in switch 1's VLANs, you must define policies on switch 1 that will so include them. In this example, an IPX policy was defined for VLAN 2 and an IP policy was defined for VLAN 3.

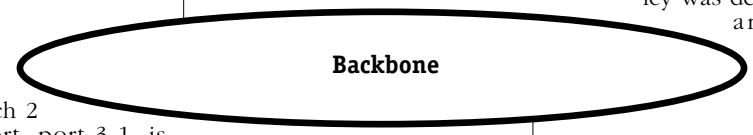

All traffic entering switch 2 from the backbone port, port 3-1, is checked against switch 2's policies for VLAN membership in the Group(s) being trunked. In this example only Group 1 is being trunked, so traffic is checked against the VLAN policies defined for Group 1's VLANs. If you want devices from switch 1 included in switch 2's VLANs, you must define policies on switch 2 that will so include them. In this example, an IPX policy was defined for VLAN 2 and an IP policy was defined for VLAN 3.

Because the backbone port, port 3-1, is assigned to VLANs 2 and 3 in switch 2 via a port policy, traffic can flow out of these VLANs onto the backbone.

OmniStack 2

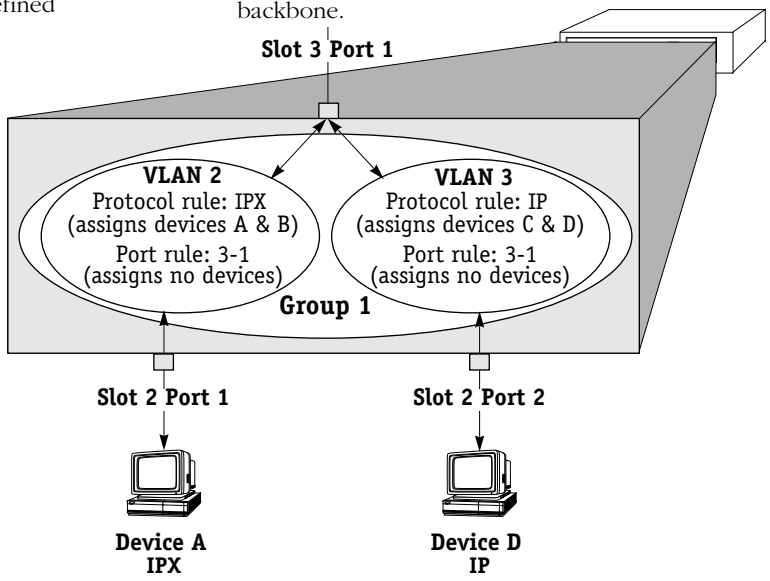

Device	Source Port	VLAN Membership
A	2-1	2
B	3-1	2
C	3-1	3
D	2-2	3

Frame Flooding in AutoTracker VLANs

Flooding occurs when a frame is received addressed to a device that is unknown to the switch or broadcast or multicast frames are received addressed to multiple users. In a typical bridged environment, the frame would be forwarded out all ports. However, this is not true with VLANs as VLANs segment the network into smaller broadcast domains. In this environment, flooding occurs as follows:

Unicast Traffic

- If the destination address of the frame is unknown buts its source address is known and the source device is a member of one or more VLANs, the frame is flooded out all ports of all VLANs in which the source device is a member. Please note the following:
  - If the source device is a member of multiple VLANs, some leakage may occur during the flooding process. Leakage may occur only among VLANs in the same Group—frames do not leak between Groups.
  - If the source device is a member of multiple VLANs and some or all of those VLANs share the same physical port, only one copy of the frame is forwarded out that port.
  - If the source device is a member of multiple VLANs that use trunking, only one copy of the frame is sent to each trunk port.
- If both the source and destination addresses of the frame are unknown, the frame is forwarded to the MPM for processing (to determine the VLAN or VLANs in which the originating device should be a member) ***and*** the frame is flooded out all ports of all VLANs in which the source port is a member.

Broadcast and Multicast Traffic

Frames are forwarded out all ports that are members of the same VLANs as the source MAC address. If the source MAC address is unknown, it is forwarded out all ports that have VLANs active on the source ports.

Routing Between AutoTracker VLANs

Devices that do not share membership in a common VLAN must use routers to communicate with one another. You can configure a virtual router port that is capable of IP and/or IPX routing for each VLAN. By enabling a router port on a VLAN, you are creating a static route entry within the switch to that VLAN. If this router port is not configured for a VLAN, then that VLAN will not be able to communicate with other VLANs unless an external router is between those VLANs. You may configure up to 16 virtual router ports within a single OmniStack. Each VLAN may contain only one router port.

Routing and the Default VLAN. You can enable routing for the default VLAN when you initially create a Group, or when you modify the Group. There are several issues about which you should be aware when enabling routing on the Default VLAN. See *Application Example 4* in Chapter 22, "AutoTracker VLAN Application Examples," for further information.

Creating AutoTracker VLANs

You create AutoTracker VLANs through the AutoTracker menu options. Creating an AutoTracker VLAN includes the following steps:

A. Enter basic information such as the name and number for the VLAN. See *Step A. Entering Basic VLAN Information* on page 20-16 for instructions on this step.

B. Define policies that define membership in the VLAN. See *Step B. Defining and Configuring VLAN Policies* on page 20-18 for instructions on this step.

C. Configure the type of routing used for communication between VLANs. In order for devices in a VLAN to communicate with devices in other VLANs, a virtual router must be configured or an external router must exist between those VLANs. See *Step C. Configuring the Virtual Router Port (Optional)* on page 20-19 for instructions on this step.

These steps are explained in detail in the sections that follow.

Step A. Entering Basic VLAN Information

1. To begin setting up the AutoTracker VLAN type **cratvl** at any prompt.

2. The following prompt displays:

   Enter the VLAN Group id for this VLAN (1):

   Enter the number for the Group to which this VLAN will belong. All VLANs belong to a Group. You can create up to 31 VLANs per Group (each Group already contains a default VLAN, VLAN #1).

3. The following prompt displays:

   Enter the VLAN Id for this VLAN (2):

   Enter the number that will identify this VLAN with the Group specified above. Up to 32 VLANs may belong to the same Group (including the default VLAN). By default the system displays the next available VLAN ID number. Press **<Enter>** to accept this default.

4. The following prompt displays:

   Enter the new VLAN's description:

   Enter a textual description that will help you identify the VLAN. For example, if you know this VLAN will be composed of only workstations using the IPX protocol, you might call the VLAN, "IPX VLAN." You may use up to 30 characters for this description.

5. The following prompt displays:

Enter the Admin Status for this vlan (Enable (e) / Disable (d):

Enter whether or not you want the Administrative Status for this VLAN to be enabled or disabled. Once enabled, the switch begins using the policies you defined. A disabled VLAN is still defined (name, number, policies intact), but the switch keeps the VLAN disabled. The enable/disable status may be changed at a later time using the **modatvl** command.

Note

A VLAN may not always be operational even when its **Admin** Status is enabled. The VLAN becomes operational as soon as a port is assigned to it. In addition, a VLAN's operation may be disabled by the switch because devices in the VLAN cease transmitting data, among other reasons.

After you enter the Administrative Status, additional prompts display that allow you to select the rules governing membership in this VLAN. Go on to the next section, *Step B. Defining and Configuring VLAN Policies* on page 20-18 to continue setting up this VLAN.

Step B. Defining and Configuring VLAN Policies

You can define AutoTracker policies by port, MAC address, protocol, network address, user definition, or port binding. You can define multiple policies for a AutoTracker VLAN if you wish. A port or device is included in a AutoTracker VLAN if it matches any one rule. For example, you can define rules based on ports, rules based on MAC address, and rules based on protocol in the same AutoTracker VLAN. However, defining multiple rules is not trivial – exercise extreme care when you do so and make sure that you understand the consequences of your definitions. In most situations, it is advisable to use one of AutoTracker's predefined rules.

Instructions for defining each AutoTracker policy type are included in Chapter 17, "Configuring Group and VLAN Policies." Follow the directions in that chapter for the policy you wish to set up.

The sections below provide directions for setting up each type of AutoTracker policy. Follow the directions for the policy you wish to set up.

1. When are done specifying AutoTracker policies the following prompt displays:

```
Configure more rules for this vlan (y/n):
```

You can set up multiple rules for the same VLAN. Enter a **Y** here if you want to set up more rules in addition to the Network Address rule specified here. If you enter **Y**, you will be prompted for the next rule that you want to set up on this VLAN. Follow the directions in the appropriate section to configure that rule.

If you enter **N**, you will receive a message, similar to the one below, indicating that the VLAN was set up.

```
VLAN 1:2 created successfully
```

You are done setting up rules for this VLAN, so you can start configuring the virtual router for this VLAN. See *Step C. Configuring the Virtual Router Port (Optional)* on page 20-19 for information on configuring a virtual router port.

Step C. Configuring the Virtual Router Port (Optional)

You can now optionally configure the virtual router port that this VLAN will use to communicate with other AutoTracker VLANs. A virtual router port for the VLAN is created within the switch. If you do not define a virtual router port for this VLAN, devices within the VLAN will only be able to communicate with devices in other VLANs through an external router.

You will have the choice of configuring IP, IPX, or both IP and IPX routing. Continue with the steps below:

1. After you finish configuring AutoTracker Policies for this VLAN, the following prompt displays:

   Enable IP (y):

   Press **<Enter>** if you want to enable IP Routing on this virtual router port. If you do not enable IP, then this VLAN will not be able to internally route IP data. If you don't want to set up the IP router port, enter **n**, press **<Enter>** and skip to Step 10.

   Note

   You may enable routing of both IP and IPX traffic on this router port. If you set up dual-protocol routing, you must fill out information for both IP and IPX parameters.

2. The following prompt displays:

   IP Address:

   Enter the IP address for this virtual router port in dotted decimal notation or hexadecimal notation (e.g., 198.206.181.10). This IP address is assigned to the virtual router port for this VLAN. After you enter the address, press **<Enter>**.

3. The following prompt displays:

   IP Subnet Mask (0xffffff00):

   The default IP subnet mask (in parentheses) is automatically derived from the VLAN's IP address class. Press **<Enter>** to select the default subnet mask or enter a new subnet mask in dotted decimal notation or hexadecimal notation and press **<Enter>**.

4. The following prompt displays:

   IP Broadcast Address (198.200.10.255):

   The default IP broadcast address (in parentheses) is automatically derived from the VLAN's IP address class. Press **<Enter>** to select the default address or enter a new IP broadcast address in dotted decimal notation or hexadecimal notation and press **<Enter>**.

5. The following prompt displays:

   Description (30 chars max):

   Enter a useful description for this virtual IP router port using alphanumeric characters. The description may be up to 30 characters long. Press **<Enter>**.

6. The following prompt displays:

   Disable routing? (n) :

   Indicate whether you want to disable routing in the VLAN. You can enable routing later through the **modvl** command.

7. The following prompt displays:

```
Enable NHRP?              (n) :
```

Indicate whether you want to enable NHRP.

8. The following prompt displays:

```
IP RIP Mode {Deaf (d),
    Silent (s),
    Active (a),
    Inactive (i)}                 (s):
```

Define the RIP mode in which the virtual router port will operate. RIP (Router Information Protocol) is a network-layer protocol that enables this VLAN to learn and advertise routes. The RIP mode can be set to one of the following:

Silent. The default setting shown in parentheses. RIP is active and receives routing information from other VLANs, but does not send out RIP updates. Other VLANs will not receive routing information concerning this VLAN and will not include the VLAN in their routing tables. Simply press **<Enter>** to select Silent mode.

Deaf. RIP is active and sends routing information to other VLANs, but does not receive RIP updates from other VLANs. This VLAN will not receive routing information from other VLANs and will not include other VLANs in its routing table. Enter **d** and press **<Enter>** to select Deaf mode.

Active. RIP is active and both sends and receives RIP updates. This VLAN will receive routing information from other VLANs and will be included in the routing tables of other VLANs. Enter **a** and press **<Enter>** to select Active mode.

Inactive. RIP is inactive and neither sends nor receives RIP updates. This VLAN will neither send nor receive routing information to/from other VLANs. Enter **i** and press **<Enter>** to select Inactive mode.

9. After you enter the RIP mode, the following prompt displays:

```
Default framing type [Ethernet II(e),
        fddi (f),
        token ring (t),
        Ethernet 802.3 SNAP (8),
        source route token ring(s)} (e):
```

Select the default framing type for the frames that will be generated by this router port and propagated over this VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in this VLAN. If this VLAN contains devices using encapsulation types other than those defined here, the MPM module must translate those frames, which slows throughput. The figure on the next page illustrates the Default Framing Type and its relation to Virtual Router Port communications.

After you enter the framing type a message displays indicating that this IP router port was created:

```
Created router port for vlan 1:3
```

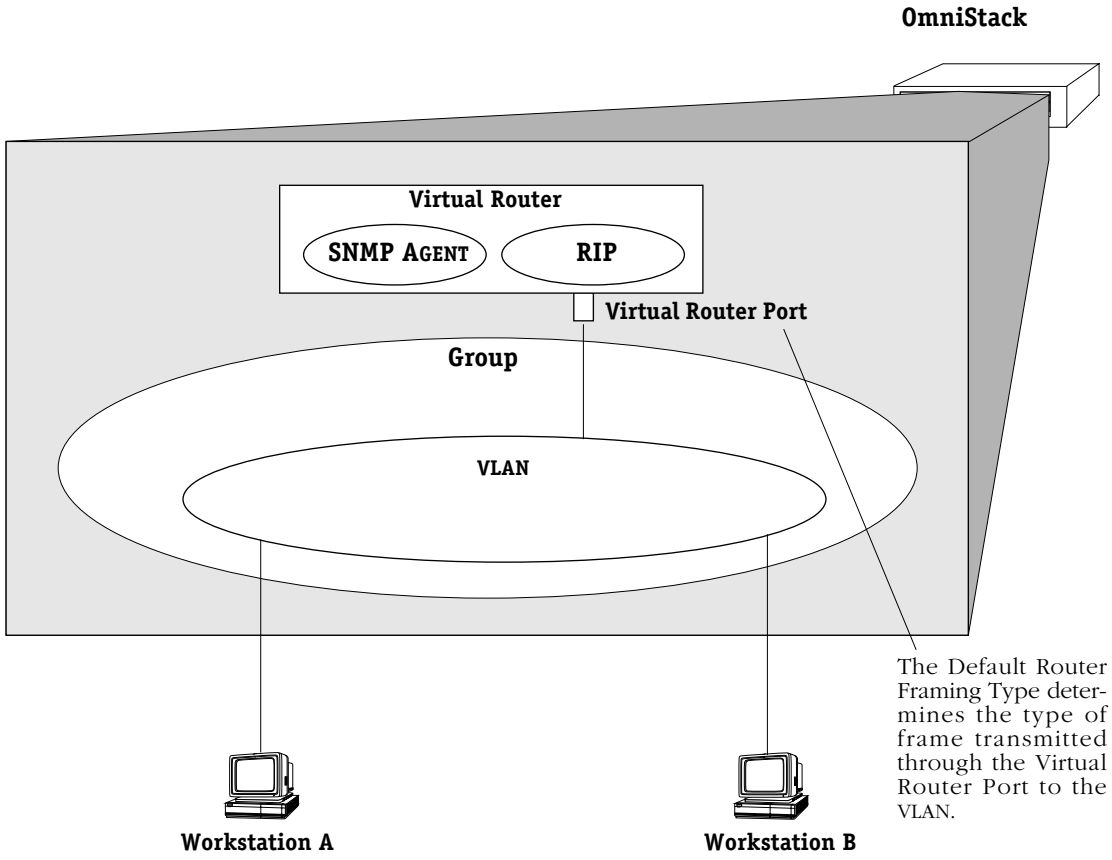

Default Framing Type and the Virtual Router Port

10. You can now configure IPX routing on this port. The following message displays:

Enable IPX? (y) :

Press **<Enter>** if you want to enable IPX Routing on this virtual router port. If you do not enable IPX, then this VLAN will not be able to internally route IPX data. You can set up a virtual router port to route both IP and IPX traffic.

If you don't want to enable IPX routing, enter **n** and press **<Enter>.** You are now done configuring this VLAN. You can monitor activity on this VLAN through other AutoTracker commands. See later section in this chapter for more information on these commands.

11. After selecting to enable IPX, the following prompt displays:

IPX Network:

Enter the IPX network address. IPX addresses consist of eight hex digits and you can enter a minimum of one hex digit in this field. If you enter less than eight hex digits, the system prefixes your entry with zeros to create eight digits.

12. The following prompt displays:

Description (30 chars max):

Enter a useful description for this virtual IPX router port using alphanumeric characters. The description may be up to 30 characters long. Press **<Enter>**.

13. After entering a description, the following prompt displays:

```
IPX RIP and SAP mode {RIP and SAP active (a)
        RIP only active (r)
        RIP and SAP inactive (i)}                    (a):
```

Select how you want the IPX protocols, RIP (router internet protocol) and SAP (service access protocol), to be configured for this VLAN. RIP is a network-layer protocol that enables this VLAN to learn routes. SAP is also a network-layer protocol that allows network services, such as print and files services, to advertise themselves. The choices are:

RIP and SAP active. The default setting. The VLAN to which this IPX router port is attached participates in both RIP and SAP updates. RIP and SAP updates are sent and received through this router port. Simply press **<Enter>** to select RIP and SAP active.

RIP only active. The VLAN to which this IPX router port is attached participates in RIP updates only. RIP updates are sent and received through this router port. Enter an **r** and press **<Enter>** to select RIP only active.

RIP and SAP inactive. The IPX router port is active, but the VLAN to which it is attached does not participate in either RIP nor SAP updates. Enter an **i** and press **<Enter>** to select RIP only active.

14. After selecting the RIP and SAP configuration, the following prompt displays the default router framing type options:

```
Default router framing type for : {

        Ethernet Media:
                Ethernet II (0),
                Ethernet 802.3 LLC (1),
                Ethernet 802.3 SNAP (2),
                Novell Ethernet 802.3 raw (3),

        FDDI Media:
                fddi SNAP (4),
                source route fddi SNAP (5),
                fddi LLC (6),
                source route fddi LLC (7),

        Token Ring Media:
                token ring SNAP (8),
                source route token ring SNAP (9),
                token ring LLC (a),
                source route token ring LLC (b) }          (0) :
```

Select the default framing type for the frames that will be generated by this router port and propagated over the VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in the VLAN. If the VLAN contains devices using encapsulation types other than those defined here, the MPM module must translate those frames, which slows throughput. See the figure, *Default Framing Type and the Virtual Router Port* on page 20-21 for an illustration of the Default Framing Type and its relation to Virtual Router Port communications.

15. If you chose a Source Routing frame format in the last step (options 5, 7, 9, or b), the an additional prompt displays:

```
Default source routing broadcast type : {
      ARE broadcasts(a), STE broadcasts(s)}          (a) :
```

Select how broadcasts will be handled for Source Routing. The choices are:

ARE broadcasts. All Routes Explorer, the default setting. Broadcasts are transmitted over every possible path on inter-connected source-routed rings. This setting maximizes the generality of the broadcast. Simply press **<Enter>** to select All Routes Explorer.

STE broadcasts. Spanning Tree Explorer. Broadcasts are transmitted only over Spanning Tree paths on inter-connected source-routed rings. This setting maximizes the efficiency of the broadcast. Enter an **s** and press **<Enter>** to select Spanning Tree Explorer.

After you enter framing type information a message displays indicating that this IPX router port was created:

```
Created router port for vlan 1:3
```

You have now completed the configuration of the virtual router port for this VLAN. You can monitor activity on this VLAN through other AutoTracker commands. See later section in this chapter for more information on these commands.

Modifying an AutoTracker VLAN

After you set up a VLAN you can modify its Admin Status, description, rules, and the Admin Status of each of the rules. You use the **modatvl** command to modify a VLAN as follows:

```
modatvl <Group Number>:<VLAN Number>
```

You must specify the Group and VLAN numbers and they must be separated by a colon. For example, to modify the VLAN 3 in Group 4, you would specify:

```
modatvl 4:3
```

After entering a valid **modatvl** command a screen similar to the sample below displays:

```
VLAN    4 :  3 is defined as:
             1.      Description      = AT VLAN 3
             2.      Admin Status     = Enabled
             3.      Rule Definition
                     Rule Num    Rule Type       Rule Status
                       1         Protocol Rule   Disabled
Available options:
             1.      Set VLAN Admin Status
             2.      Set VLAN Description
             3.      Add more rules
             4.      Delete a rule
             5.      Set rule Admin Status
             6.      Quit
Option =
```

The first half of the display shows the current configuration of this VLAN. For example, this sample shows a VLAN 3 in Group 4 with a description, "AT VLAN 3." The VLAN is Enabled and a Protocol Rule has been set up, but this rule has not been enabled.

The second half of the displays a list of the VLAN attributes you can modify. You can modify basic information such as the Admin Status and Description. You can also add rules, delete rules, and enable or disable the rule. To modify an attribute, enter the number next to the option you want to modify and press **<Enter>**.

The following sections describe each of the six Available Options for the **modatvl** command.

Changing a VLAN's Admin Status

1. At the **Option=** prompt enter a **1** and press **<Enter>**.
2. The following prompt displays:

   ```
   Set Admin Status to ((e)nable/(d)isable):
   ```

   Type an **E** to enable the VLAN or a **D** to disable it. An enabled VLAN starts using policies to direct data flow. A disabled VLAN is saved, but can not become active.

   The system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Changing a VLAN's Description

1. At the **Option=** prompt enter a **2** and press **<Enter>**.
2. The following prompt displays:

   Enter a new description:

   Type in the revised description for this VLAN. The description can be up to 30 characters long. Press **<Enter>** when you have completed the new description.

   The system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Adding More Policies for This VLAN

1. At the **Option=** prompt enter a **3** and press **<Enter>**.
2. The following menu displays:

   Select rule type:
   1. Port Rule
   2. MAC Address Rule
   3. Protocol Rule
   4. Network Address Rule
   5. User Defined Rule
   6. Binding Rule
   7. DHCP PORT Rule
   8. DHCP MAC Rule

   Enter rule type (1):

   This is the same menu used by the **cratvl** command. This menu has eight options, some of which contain multiple branching options. This menu is documented fully in Chapter 17, "Configuring Group and VLAN Policies." Please consult that chapter for information on this menu.

   When have entered all new rule types, the system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Deleting A Policy for This VLAN

1. At the **Option=** prompt enter a **4** and press **<Enter>**.
2. The following menu displays:

   Enter rule number to delete:

   The rule number is listed with other information on the VLAN just after you entered the **modatvl** command. Find the number corresponding to the rule you want to delete and enter it at this prompt and press **<Enter>**. The rule is deleted and the system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Changing the Admin Status for a VLAN Policy

1. At the **Option=** prompt enter a **5** and press **<Enter>**.
2. The following menu displays:

   Enter rule number:

   The rule number is listed with other information on the VLAN just after you entered the **modatvl** command. Find the number corresponding to the rule you want to change and enter it at this prompt and press **<Enter>**.
3. The following menu displays:

   Set Rule Admin Status to ((e)nable/(d)isable):

   Type an **E** to enable this rule or a **D** to disable it. If the rule is enabled, the VLAN will start using the rule criteria to segment data traffic.

   The system returns to the **Available Options** menu. You can modify more attributes for this VLAN, or quit modifying the VLAN by typing a **6**.

Deleting an AutoTracker VLAN

You can delete an AutoTracker VLAN. When you delete a VLAN, traffic is no longer filtered according to the VLAN's policies. Follow these steps to delete a VLAN.

1. Enter **rmatvl** followed by the Group number, a colon (:), and the VLAN number that you want to delete. For example to delete VLAN 2 in Group 3, you would enter:

   rmmcvl 3:2
2. The following prompt displays:

   Delete VLAN 3:2 ? (n):

   Enter a **Y** and press **<Enter>** to complete the deletion of the VLAN. A message display confirming the deletion.

   VLAN 3:2 deleted

Viewing AutoTracker VLANs

You can view the current status of all AutoTracker VLANs in the switch using the **atvl** command. Enter atvl and a table similar to the following displays.

VLAN Group :	VLAN Id	VLAN Description	Admin Status	Operational Status
3:	5	VLAN 5	Enabled	Active
3:	11	VLAN 11	Enabled	Inactive
3:	12	VLAN 12	Enabled	Inactive
3:	22	VLAN 22	Enabled	Active
3:	23	VLAN 23	Enabled	Active
3:	24	VLAN 24	Enabled	Inactive
3:	25	VLAN 25	Enabled	Inactive
3:	26	VLAN 26	Enabled	Inactive
3:	27	VLAN 27	Enabled	Inactive
3:	31	VLAN 31	Enabled	Inactive
3:	32	VLAN 32	Enabled	Inactive

VLAN Group. The Group to which this AutoTracker VLAN is assigned. The Group is specified when first creating an AutoTracker VLAN.

VLAN ID. An identification number that you assigned when you created this VLAN.

VLAN Description. A textual description that you entered to describe a VLAN when you created or modified it through **cratvl** or **modatvl**. This description is limited to 30 characters.

Admin Status. The Administrative Status for the VLAN may be enabled or disabled. You enable or disable the Administrative Status for a VLAN when you create or modify it. If the VLAN is enabled, the switch will use the policies you configured to filter traffic to the devices in this VLAN. If you disable the rule, then policies will not be used, but the parameters you set up for the VLAN will be saved.

Oper Status. The VLAN is shown as **Active** or **Inactive**. In order for an enabled VLAN to become "active" it must be able to assign a switch port to the VLAN. If the port rule is used for a VLAN, then the VLAN automatically becomes active. If any other rule is used (MAC address, protocol, etc.), then a frame matching the VLAN rule must first be received by a switch port before the VLAN is active. So, an Active VLAN requires the following:

- Admin Status must be enabled.
- A port must be assigned to the VLAN through either a port-based rule or by a device transmitting data that matches the VLAN policy.

Viewing Policy Configurations

Typing **viatrl** brings up the Policy Configuration Table, which shows the policies defined for the VLAN specified.

VLAN Group :	VLAN Id	Rule Num	Rule Type	Rule Status	Rule Definition
3:	5	1	PORT RULE	Disabled	2/7/Brg/1
3:	11	1	NET ADDR RULE	Enabled	IPX Addr = 11223344
					IPX Encapsulation = Ethernet
3:	12	1	NET ADDR RULE	Enabled	DECNET Area = 13579
3:	22	1	PORT RULE	Enabled	2/7/Brg/1
3:	23	1	PORT RULE	Enabled	2/7/Brg/1
3:	24	1	MAC RULE	Enabled	082008:003002
					082009:803728
3:	25	1	PROTOCOL RULE	Enabled	Protocol = IP
3:	26	1	NET ADDR RULE	Enabled	IP Addr = 131.1.2.3
					IP Mask = 255.255.0.0
3:	27	1	USER RULE	Enabled	Offset = 64
					Length = 2
					Value = FFFF
					Mask = FFFF
3:	31	1	PROTOCOL RULE	Enabled	Protocol = IP
3:	32	1	NET ADDR RULE	Enabled	IPX Addr = 00000001
					IPX Encapsulation = Ethernet

VLAN Group. The Group to which this AutoTracker VLAN is assigned. The Group number is specified when first creating the VLAN.

VLAN ID. An identification number that you assigned when you created this virtual LAN.

Rule Num. The number of the policy within the VLAN definition. Each rule defined for a VLAN is numbered sequentially in the order of creation. The rule number is needed when you want to modify or delete a rule definition.

Rule Type. The type of VLAN policy. The Rule Type can be a port policy (PORT RULE), MAC Address policy (MAC RULE), network address policy (NET ADDR RULE), Protocol policy (PROTOCOL RULE), or a user-defined policy (USER RULE). You set up VLAN policies when you create or modify the VLAN.

Rule Status. Indicates whether the rule for this row is Enabled or Disabled. If the rule is enabled, then the VLAN is using the rule definition to determine VLAN membership. If Disabled, then the VLAN is not using this rule to determine membership. Note that this Rule Status is different from the Admin Status for the VLAN since it controls only this specific rule within this specific VLAN. You can enable or disable the rule using the **modatvl** command.

Rule Definition. Details of this rule. For a Port Rule, this column lists the virtual interface for the Port included in the VLAN as

<slot>/<port>/<service>/<instance>

For example, the port defined for the first row in the table applies to the first bridge instance on port 7 on the module in slot 2 of the switch. For a MAC address rule, this column lists the MAC address for the device in the VLAN. For a Network Address Rule, the column will list the address (IP or IPX) and the IP Mask (IP) or the Encapsulation type (IPX). For a Protocol policy, the column list the protocol used to determine membership. And in a User-Defined rule, the offset, length, value, and mask are listed.

Viewing Virtual Ports' VLAN Membership

You can view the VLAN membership of each virtual interface in the switch. For physical LAN ports, the virtual interface is the same as a virtual port. However, when multiple services are set up for a physical port, then each service has a virtual port.

Type **vivl** and a Virtual Interface Table displays similar to the one that follows. You can also specify just the slot and port number to narrow the range of ports displayed.

Virtual Interface VLAN Membership

Slot/Intf/Service/Instance	Group	Member of VLAN#
1 /1 /Rtr /1	1	1
1 /1 /Rtr /2	3	1
1 /1 /Rtr /3	3	23
1 /1 /Rtr /4	3	24
1 /1 /Rtr /5	3	25
1 /1 /Rtr /6	3	5
2 /1 /Brg /1	1	1
2 /2 /Brg /1	1	1
2 /3 /Brg /1	1	1
2 /4 /Brg /1	1	1
2 /5 /Brg /1	1	1
2 /6 /Brg /1	1	1
2 /7 /Brg /1	1	1 22
2 /8 /Brg /1	1	1
3 /1 /Brg /1	1	1
4 /1 /Brg /1	1	1
4 /2 /Brg /1	1	1
4 /3 /Brg /1	1	1
4 /4 /Brg /1	1	1
4 /5 /Brg /1	1	1
4 /6 /Brg /1	1	1
5 /1 /Brg /1	1	1

Slot/Intf/Service/Instance. Specifies the virtual interface for which AutoTracker VLAN information will be displayed. The **Slot** is the physical slot location to which the virtual interface maps. The **Intf** is the physical port to which the virtual interface maps. The **Service** is the service type for this interface. The service type may be a Router (**Rtr**), Bridge (**Brg**), Classical IP (**CIP**), FDDI Trunk (**Trk**), or an 802.10 Trunk (**T10**). **Instance** is the specific instance of this service type. These different instances are identified numerically. The first instance of a service type belonging to a physical port is identified as 1, the second instance is identified as 2, etc.

Group. The Group to which this virtual interface is assigned. The Group is specified when first creating an AutoTracker VLAN.

Member of VLAN #. The AutoTracker VLANs to which this virtual interface belongs. An interface may belong to more than one VLAN. For example, a port may contain devices using the IP Protocol and could match the Port policy of one AutoTracker VLAN and the Protocol policy of another AutoTracker VLAN. Also, physical ports always remain members of the default VLAN #1.

View VLAN Membership of MAC Devices

The **fwtvl** command displays a table of learned MAC addresses and the VLAN membership of those MAC addresses. Follow these steps to view this table.

1. Enter **fwtvl**.
2. The following prompt displays:

   Enter Slot/Interface (return for all ports) :

   Enter the slot and port for which you want to view MAC Address/VLAN information. You can also press **<Enter>** to view information on all ports in the switch.
3. The following message and prompt displays:

   Total number of MAC addresses learned for Group 1: 4
   Maximum number of entries to display [20]) :

   The top line displays the number of MAC addresses learned on this switch. This number indicates the potential number of entries you can display in the Learned MAC Address Table. The second line allows you to indicate how many of these MAC addresses you want to display. Enter the number of MAC entries you want to display or press **<Enter>** to select the default in brackets [20].
4. The Learned MAC Address/VLAN Membership Table displays as follows:

MAC Address	Slot/Intf/Service/Instance	AT VLAN Membership
0020DA:05F623	4/ /1 /Brg 1	1
0020DA:021533	4/ /1 /Brg 1	1
0020DA:0205B3	4/ /1 /Brg 1	1
0020DA:06BAD3	4/ /1 /Brg 1	1
0020DA:05F610	4/ /1 /Brg 1	1

MAC Address. The MAC address for which virtual interface and VLAN membership information will be displayed.

Slot/Intf/Service/Instance. Specifies the virtual port for which AutoTracker VLAN information will be displayed. The **Slot** is the physical slot location to which the MAC address maps. The **Intf** is the physical port to which the MAC address maps. The **Service** is the service type for this MAC address. The service type may be a Router (**Rtr**), Bridge (**Brg**), Classical IP (**CIP**), FDDI Trunk (**Trk**), or an 802.10 Trunk (**T10**). **Instance** is the specific instance of this service type. These different instances are identified numerically. The first instance of a service type belonging to a physical port is identified as 1, the second instance is identified as 2, etc.

AT VLAN Membership. The AutoTracker VLANs to which this MAC Address belongs. An MAC address may belong to more than one VLAN. For example, let's say a MAC device runs on an IPX network. It could be included in a MAC Address policy for one AutoTracker VLAN and the IPX Protocol Policy of another VLAN.

Creating a VLAN for Banyan Vines Traffic

Banyan Vines uses a fixed encapsulation for each network interface. For this reason, it is straightforward to create a VLAN for Banyan Vines traffic. For Ethernet traffic, Banyan Vines uses Ethernet II encapsulation. Follow these steps to create a Banyan Vines VLAN:

1. Type **cratvl** at any prompt.
2. The following prompt displays:

   Enter the VLAN Group id for this VLAN (1):

   Enter the number for the Group to which this Banyan Vines VLAN will belong.
3. The following prompt displays:

   Enter the VLAN Id for this VLAN (2):

   Enter the number that will identify this VLAN within the Group specified above. By default the system displays the next available VLAN ID number. Press **<Enter>** to accept this default.
4. The following prompt displays:

   Enter the new VLAN's description:

   Enter a textual description that will help you identify the VLAN. For example, you might call the VLAN, "Banyan Vines VLAN." You may use up to 30 characters for this description.
5. The following prompt displays:

   Enter the Admin Status for this vlan (Enable (e) / Disable (d):

   Enter whether or not you want the Administrative Status for this VLAN to be enabled or disabled. Once enabled, the switch begins using the policies you defined. A disabled VLAN is still defined (name, number, policies intact), but the switch keeps the VLAN disabled. The enable/disable status may be changed at a later time using the **modatvl** command.
6. The following menu displays:

   Select rule type:
   1. Port Rule
   2. MAC Address Rule
   3. Protocol Rule
   4. Network Address Rule
   5. User Defined Rule
   6. Binding Rule
   7. DHCP PORT Rule
   8. DHCP PORT Rule

   Enter rule type (1):

   Press **3** and press **<Enter>**.
7. The following prompt displays:

   Set Rule Admin Status to ((e)nable/(d)isable):

   Type **e** to enable this rule. When enabled, the VLAN will begin using the rule to determine membership of devices.

8. The following prompt displays:

```
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type
6. Protocol specified by DSAP and SSAP
7. Protocol specified by SNAP

Enter protocol type (1):
```

Enter a **5** to define a protocol by ether-type and press **<Return>**.

9. The following prompt displays:

```
Enter the Ether-type value in hex:
```

10. Enter **0bad** as the Ether-type value for Ethernet II encapsulation.

11. The following prompt displays:

```
Configure more rules for this vlan (y/n):
```

Enter a **Y**. You still need to set up rules for LLC and SNAP traffic.

12. The following prompt displays:

```
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP PORT Rule

Enter rule type (1):
```

Press **3** and press **<Enter>**.

13. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Type **e** to enable this rule.

14. The following prompt displays:

```
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type
6. Protocol specified by DSAP and SSAP
7. Protocol specified by SNAP

Enter protocol type (1):
```

Enter a **6** to define a protocol by DSAP and SSAP and press **<Return>**.

15. The following prompt displays

```
Enter the DSAP value in hex:
```

Enter **bc** as the destination service access protocol (DSAP) value and press **<Enter>**.

16. The following prompt displays:

```
Enter the SSAP value in hex:
```

Again, enter **bc** as the source service access protocol (SSAP) value and press **<Enter>**.

17. The following prompt displays:

```
Configure more rules for this vlan (y/n):
```

Enter a **Y**. You still need to set up a rule for SNAP traffic.

18. The following prompt displays:

```
Select rule type:
1. Port Rule
2. MAC Address Rule
3. Protocol Rule
4. Network Address Rule
5. User Defined Rule
6. Binding Rule
7. DHCP PORT Rule
8. DHCP PORT Rule

Enter rule type (1):
```

Press **3** and press **<Enter>**.

19. The following prompt displays:

```
Set Rule Admin Status to ((e)nable/(d)isable):
```

Type **e** to enable this rule.

20. The following prompt displays:

```
Select Protocol:
1. IP
2. IPX
3. DECNET
4. APPLETALK
5. Protocol specified by ether-type
6. Protocol specified by DSAP and SSAP
7. Protocol specified by SNAP

Enter protocol type (1):
```

Enter a **7** to define a protocol by SNAP and press **<Return>**.

21. The following prompt displays:

```
Enter the SNAP value in hex
```

Enter 00000080c4 as the desired SNAP value and press **<Return>**.

22. The following prompt displays:

Configure more rules for this vlan (y/n):

Enter an **N**. You are done setting up rules for this VLAN. A prompt similar to the following displays:

VLAN 1:2 created successfully

23. The following prompt displays:

Enable IP (y):

Enter an **N**.

24. The following prompt displays:

Enable IPX (y):

Enter an **N**. The Banyan Vines traffic VLAN is complete.