

22 AutoTracker VLAN Application Examples

This chapter provides specific examples of AutoTracker VLANs in various network configurations. These examples illustrate basic concepts about AutoTracker and highlight issues that can arise when AutoTracker is used in different network situations.

- *Application Example 1* illustrates a network organized according to logical policies and explains the benefits of a logical network organization.
- *Application Example 2* explains unique characteristics of IPX networks that must be considered when using AutoTracker IPX network address VLANs.
- *Application Example 3* highlights an issue concerning translated frames and AutoTracker IPX network address VLANs.
- *Application Example 4* explains how routing works generally in IPX networks and explains how to avoid an exception condition in which AutoTracker can affect the behavior of an IPX-routed network.
- *Application Example 5* explains why a port-based policy may be required for a VLAN – in addition to any other policies defined for that VLAN – to establish communications in some network situations, such as traversing a backbone.

Application Example 1

VLANs Based on Logical Policies

Example 1 shows a network organized logically. The network is organized according to IP networks, but this organization is achieved through the application of logical policies rather than physical segmentation. The use of logical policies provides the flexibility of moving IP users from segment to segment and preserving their original VLAN membership – without reconfiguring AutoTracker or the workstations.

Group and VLAN Membership

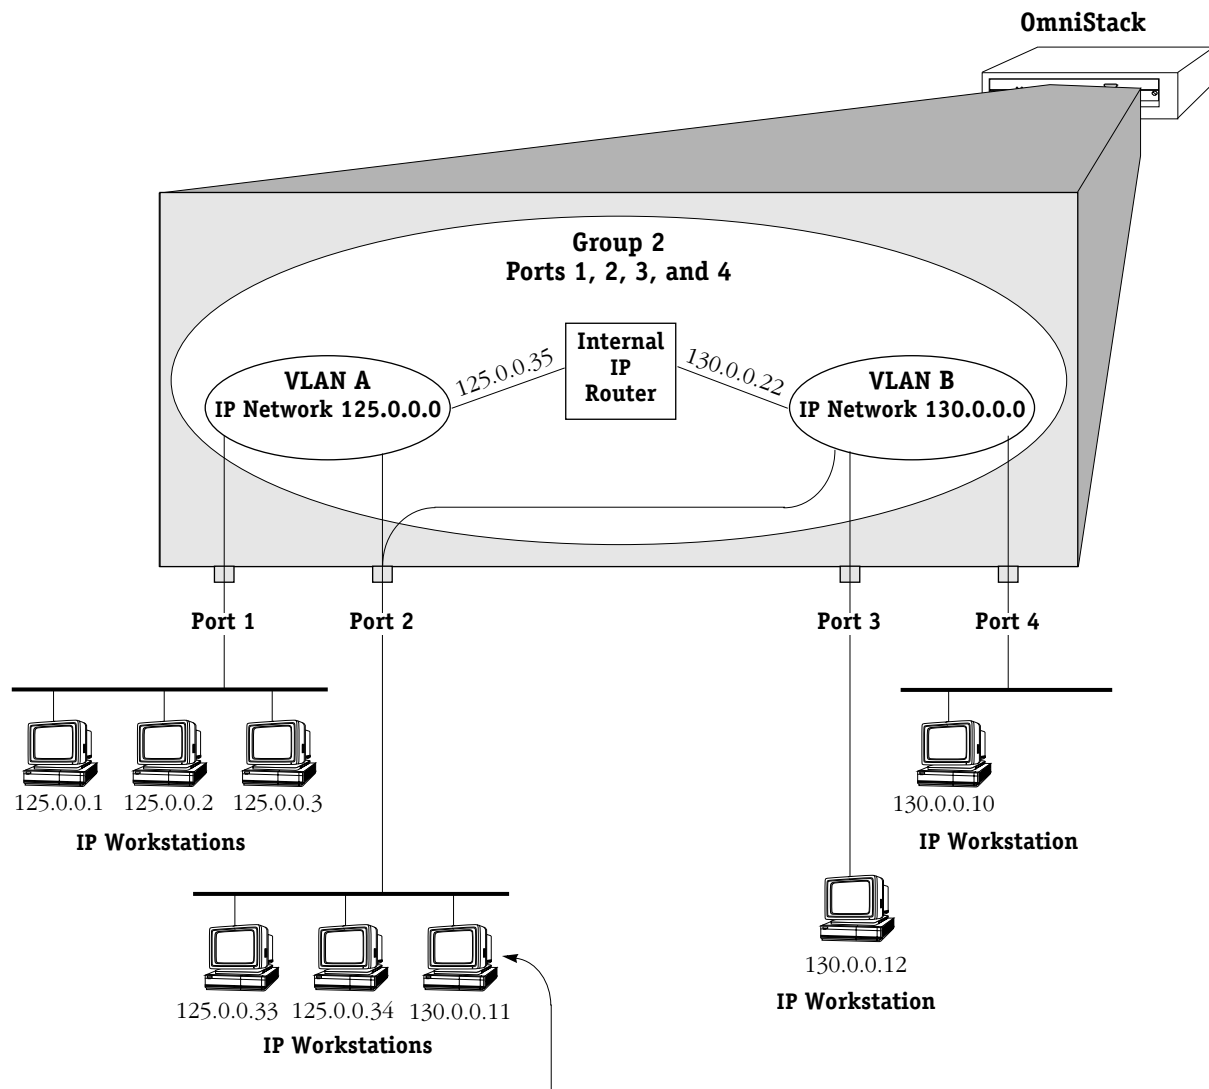
The network shown in Example 1 contains one Group – Group 2 – that consists of ports 1, 2, 3, and 4. Note that a Group defines a physical area – a set of ports – within the network. When VLANs with logical policies are created within a Group, the logical policies are applied to traffic received from all ports within the Group – but not to traffic from ports outside the Group – to determine if any source device should be a VLAN member.

As shown on the facing page, two VLANs were created within Group 2, each with a logically-based Network Address policy. The Network Address policy for VLAN A defines IP network 125.0.0.0 and the Network Address policy for VLAN B defines IP network 130.0.0.0. All traffic received on ports 1, 2, 3, and 4 will be checked for possible membership in these two VLANs.

Routing was enabled on both VLAN A and VLAN B so that traffic can move between the two VLANs, as is shown in this example by the presence of the internal IP router.

Benefits

This network configuration shown in this example provides flexibility. As explained on the following page, this logical network organization enables the Network Manager to move IP users between segments while preserving their original VLAN membership – without reconfiguring AutoTracker or the workstations.



Workstation 130.0.0.11 has been moved from the segment connected to port 4 to the segment connected to port 2. When workstation 130.0.0.11 transmits its first frame from its new location, the switch automatically places it into its original VLAN, VLAN B, because VLAN B has a network address rule that places all devices with network address 130.0.0.0 into VLAN B.

Both VLAN A and VLAN B are now active on port 2. In addition, VLAN B is now active on multiple ports – ports 2, 3, and 4. However, this does not cause confusion.

As an example, if workstation 125.0.0.1 (in VLAN A) wants to talk to workstation 130.0.0.11 (in VLAN B), workstation 125.0.0.1 ARPs for workstation 130.0.0.11's MAC address. The address returned is that of workstation 125.0.0.1's default gateway, which is VLAN A's internal IP router, 125.0.0.35. Workstation 125.0.0.1 transmits its frame to this address and the internal IP router routes the frame to VLAN B.

When VLAN B's internal IP router receives the frame addressed to workstation 130.0.0.11, it ARPs for workstation 130.0.0.11's MAC address if it does not already know it. The switch's filtering database identifies the port through which this MAC address can be reached. The frame sent by workstation 125.0.0.1 to workstation 130.0.0.11 is correctly transmitted to port 2.

Application Example 2

VLANs in IPX Networks

Example 2 illustrates the use of AutoTracker VLANs in IPX networks – specifically, VLANs based on IPX network address rules. IPX networks have unique characteristics that must be considered when configuring VLANs based on network address rules.

Encapsulation Type in IPX Networks

The encapsulation type a MAC station uses is very important in IPX networks, because a close relationship exists between encapsulation type and IPX network number. In IPX networks, a network number and an encapsulation type are configured for each segment. When two IPX servers share the same LAN segment, they must have the same network number and the same encapsulation type in order to communicate. In addition, only clients and servers that use the same encapsulation type can communicate. (The OmniStack removes this restriction somewhat through MAC-layer translations, which will not be discussed at this time.)

In summary, network number and encapsulation type define a broadcast domain in an IPX network that is analogous to a LAN – or a VLAN. (Remember that VLANs have the same characteristics as LANs, with the exception that VLANs can span multiple segments as LANs cannot.)

An encapsulation type is configured within each IPX client prior to bootup on the network. An IPX client acquires its network number dynamically from an IPX server (or from an intervening router) using a “Get_Nearest_Server” mechanism. Upon bootup, each client sends a query seeking the nearest server that uses the same encapsulation type as the client. Only those servers using the same encapsulation type respond to the query. (An intervening router can also respond to the query: routers traditionally interconnect LAN segments and can use different encapsulation types for different networks.) This means that IPX clients do not know their network numbers at bootup, but rather acquire their network numbers after they have communicated with IPX servers or with an intervening router.

VLAN Assignment in IPX Networks

The close relationship between encapsulation type and network number in IPX networks is the main reason AutoTracker’s IPX network address policy requires you to specify both a network number and an encapsulation type. The OmniStack assigns devices to IPX network address VLANs as follows:

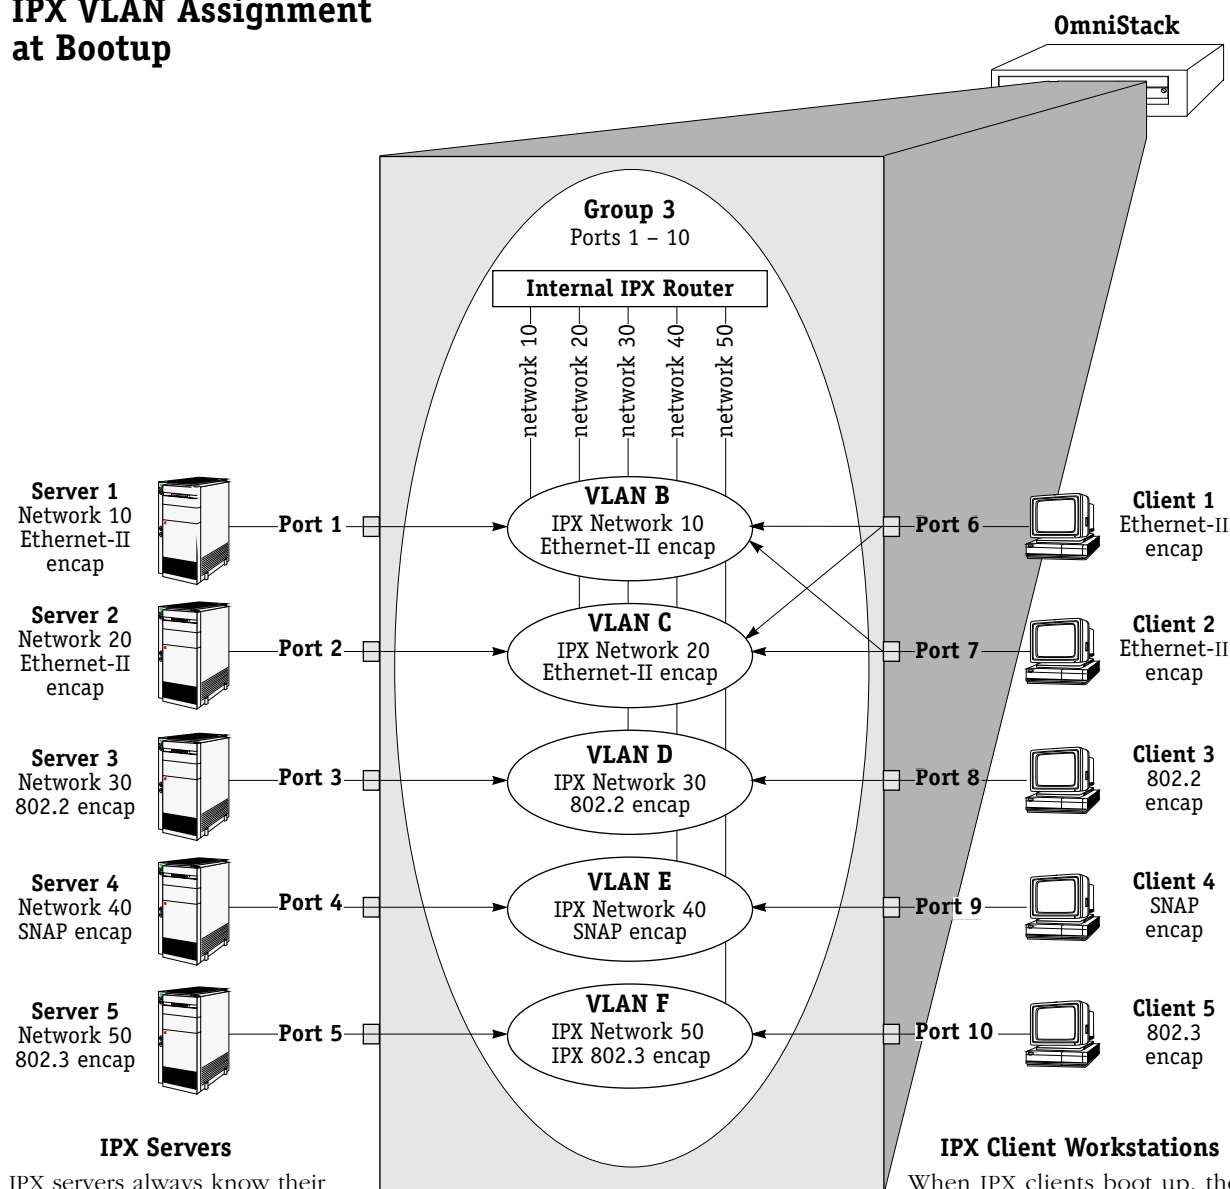
- **IPX servers.** Frames from an IPX server always contain information on the server’s network number, so the OmniStack can always assign IPX servers to the correct VLAN based on the server’s network number.
- **IPX clients.** As explained previously, IPX clients do not know their network number at bootup and so cannot, initially, be assigned to VLANs based on their network number. For this reason the OmniStack initially assigns clients to IPX network address VLANs based on their encapsulation type. An example of this is shown on the facing page. Once an IPX client communicates with a server or an intervening router, learns its network number and begins transmitting frames with that number, it is removed from all previously-assigned IPX network address VLANs (but not from VLANs of other policy types) and placed into the correct IPX network address VLAN according to network number.

So How Do I Avoid Conflicts?

As an example, IPX defines four different types of Ethernet encapsulation: Ethernet-II, 802.2, SNAP, and IPX 802.3 (also referred to as “raw”). So, what do you do to avoid conflicts when you have more than four servers and they use different encapsulation types? The solution is to put each server into a different VLAN, as shown in the example on the facing page.

continued ...

IPX VLAN Assignment at Bootup



IPX servers always know their network number, so IPX servers are assigned to VLANs according to network number.

IPX Client	VLAN Membership
Client 1	both B & C initially, then either B or C when network number is known
Client 2	both B & C initially, then either B or C when network number is known
Client 3	D
Client 4	E
Client 5	F
Please note that all ports in Group 3 are also members of Group 3's default VLAN #1.	

When IPX clients boot up, their encapsulation types are known but their network numbers are not. Therefore, IPX clients are initially assigned to VLANs according to encapsulation type. This is the reason Clients 1 and 2 (which use Ethernet-II encapsulation) are assigned to VLANs B and C (which both specify Ethernet-II encapsulation).

Once an IPX client communicates with a server or an intervening router, learns its network number and begins transmitting frames with that number, it is removed from all previously-assigned IPX VLANs and placed into a single IPX VLAN according to network number. Client 1 and Client 2 will be reassigned to either VLAN B or VLAN C when their respective network numbers are known.

In this example one Group was created – Group 3 – that includes all ports to which IPX servers and clients are connected. Within this Group five VLANs were created, one for each server:



When the OmniStack receives frames from the five servers, each server is assigned to the appropriate VLAN and no conflict occurs. IPX routing is enabled for each VLAN – with appropriate framing specified – so that traffic can route between the VLANs.

When a client workstation boots up and queries for a server, the OmniStack assigns the client to the appropriate VLAN(s) based on encapsulation type. If the client uses 802.2 encapsulation, SNAP encapsulation, or IPX 802.3 encapsulation, VLAN assignment is simple: the client is assigned to VLAN D (802.2 encapsulation), VLAN E (SNAP encapsulation), or VLAN F (IPX 802.3 encapsulation), respectively.

However, when a client workstation using Ethernet-II encapsulation boots up and queries for a server, the OmniStack initially assigns the client to both VLAN B and VLAN C, since both of these VLANs specify Ethernet-II encapsulation. However, the OmniStack recognizes that the client's frame is a "Get_Nearest_Server" query and remembers that the client is in search of its network number. While the client remains in this transitional state, it remains assigned to all VLANs that specify Ethernet-II encapsulation. Once the client has received response from a server or servers or from an intervening router, the client selects its network number and begins transmitting frames with the network number embedded. The OmniStack detects these frames, removes the client from all previously-assigned IPX network address VLANs (but not from VLANs of other policy types) and assigns it to the proper IPX network address VLAN according to network number.

Please Take Note

IPX clients often are not particular about the server to which they attach. However, clients can select a preferred server if the **/PS** (preferred server name) option is included in their start-up script.

Why is this Solution Recommended?

As as been explained, isolating each IPX server in its own IPX network address VLAN is the recommended way to avoid conflicts. No problems occur if a client receives broadcast and multicast traffic from multiple servers, especially for the brief period that the client remains in a transitional state in search of a server.

Problems do occur if two servers with different network numbers and the same encapsulation type are members of the same VLAN, because each server will detect the other's frames, notice conflicting network numbers for the same VLAN, and respond with a router configuration error. For this reason it is not advisable to create four VLANs based on IPX network address policies within the same Group, each configured for one of the four encapsulation types. It is important to isolate the servers, but it is not important to isolate the clients – at least immediately.

While it is not important to isolate IPX clients immediately at bootup, it is desirable to isolate them as soon as possible. Isolating clients – rather than letting them remain in multiple VLANs that specify the same encapsulation type – increases efficiency and reduces broadcast and multicast traffic in the network. If a client remains in multiple VLANs that specify the same encapsulation type, the client receives all broadcast and multicast traffic from each server using that encapsulation type, even though the client only communicates with the server that shares its network number. In addition, when a VLAN is extended across a WAN backbone, it is wasteful and inefficient to transmit unnecessary frames across the WAN. For these reasons, as soon as a client learns its network number and begins transmitting frames with that number, the OmniStack removes the client from all previously-assigned IPX network address VLANs and assigns it to a single IPX VLAN according to network number.

Application Example 3

IPX Network Address VLANs and Translated Frames

Application Example 3 shows two IPX networks connected over a bridged FDDI ring spanning two OmniStacks. VLAN B exists in both switches and specifies an IPX network address policy of network number 100 and Ethernet-II encapsulation.

The Problem

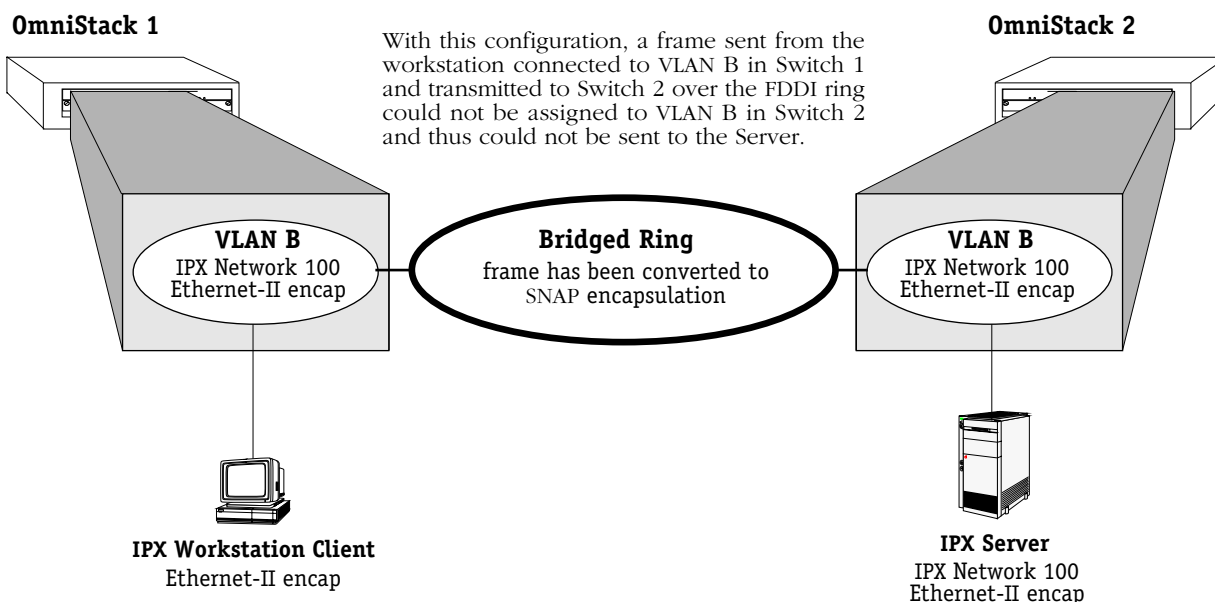
In the figure below, when the IPX client connected to Switch 1 boots up and sends a frame seeking a server, Switch 1 assigns the frame to VLAN B (since VLAN B specifies Ethernet-II encapsulation) and then converts the frame to SNAP encapsulation so that the frame can traverse the FDDI ring. When the frame arrives at Switch 2, the network number is not available (since, as previously explained, IPX clients do not know their network number at bootup) and the frame's encapsulation is no longer Ethernet-II – it is now SNAP. Because the IPX network address policy selects VLAN members according to network number and encapsulation, Switch 2 cannot assign the frame to VLAN B and send it to the IPX server.

The Solution

The solution for this problem is to specify a second encapsulation type for VLAN B in addition to Ethernet-II – for example, SNAP encapsulation. If VLAN B specifies Ethernet-II or SNAP encapsulation, the frame will match the network address policy for VLAN B when it arrives at Switch 2 and can thus be assigned to VLAN B and sent to the server. Note that the second encapsulation type must be specified for VLAN B in both Switches, to accommodate data transmission in either direction.

Please Take Note

This situation can occur whenever translations occur, such as with bridged rings. If you are using trunking you do not need to specify a second encapsulation policy for IPX network address VLANs, because trunked frames are not translated.



Application Example 4

Routing in IPX Networks

How Routing Works Generally

AutoTracker “activates” a VLAN – and its internal router interface – when the first port is assigned to the VLAN. If a VLAN has a port policy, AutoTracker assigns the specified port(s) and activates the VLAN immediately. If a VLAN has a logical policy, AutoTracker assigns the first port to the VLAN when a frame is received from a source device that matches the VLAN’s policy. When such a frame is received, the source device – and the port to which that device is connected – are assigned to the VLAN and the VLAN is activated.

Until a port is assigned to a VLAN, that VLAN is maintained in an inactive state and its internal router port is inactive – even if routing was enabled by the user. Use of a VLAN’s routing service is “on-demand” and AutoTracker does not enable routing until a port is present that might require it. When AutoTracker assigns the first port to a particular VLAN, it activates that VLAN and its routing service (as long as routing was enabled by the user).

Once AutoTracker has established devices’ VLAN assignments and activated the appropriate VLAN routing services, it does not participate in the routing process. Routing works correctly as long as the policies of the IPX protocol were followed – with the exception below.

The Exception

There is one scenario in which AutoTracker affects the behavior of an IPX-routed network. This situation occurs when an IPX server is a member of any VLAN with IPX network address policies **and** IPX routing is enabled on the Group’s default VLAN #1. An exception condition arises in this situation because all ports in a Group are always members of that Group’s default VLAN #1 in addition to any other VLANs of which they are members. As a result, default VLAN #1 is always active.

The figure on the facing page illustrates this problem situation. In this figure, three VLANs within Group 2 – one of which is default VLAN #1 – have IPX routing enabled, as indicated by the presence of the internal IPX router. VLANs 2 and 3 both have IPX network address policies. When IPX Server A is connected to the OmniStack on port 1, the Server is assigned to VLAN 2 (per the network address policy) and port 1 becomes a member of VLAN 2. When IPX Server B is connected to the OmniStack on port 2, the Server is assigned to VLAN 3 (per the network address policy) and port 2 becomes a member of VLAN 3. However, ports 1 and 2 are also members of the Group’s default VLAN #1, so port 1 is now a member of VLAN 1 and VLAN 2 and port 2 is now a member of VLAN 1 and VLAN 3.

When IPX Server A sends broadcasts, they are restricted to VLAN 2 because of the network address policies. When IPX Server B sends broadcasts, they are restricted to VLAN 3, also because of the network address policies. However, when the internal IPX router sends out broadcasts on VLAN 1 the broadcasts are flooded out all ports in the Group, because all ports in the Group are, by default, members of VLAN 1. IPX Server A responds to this with a router configuration error because it is receiving broadcasts on VLAN 1 when it should only receive them on VLAN 2. IPX Server B also responds with a router configuration error because it is receiving broadcasts on VLAN 1 when it should only receive them on VLAN 3.

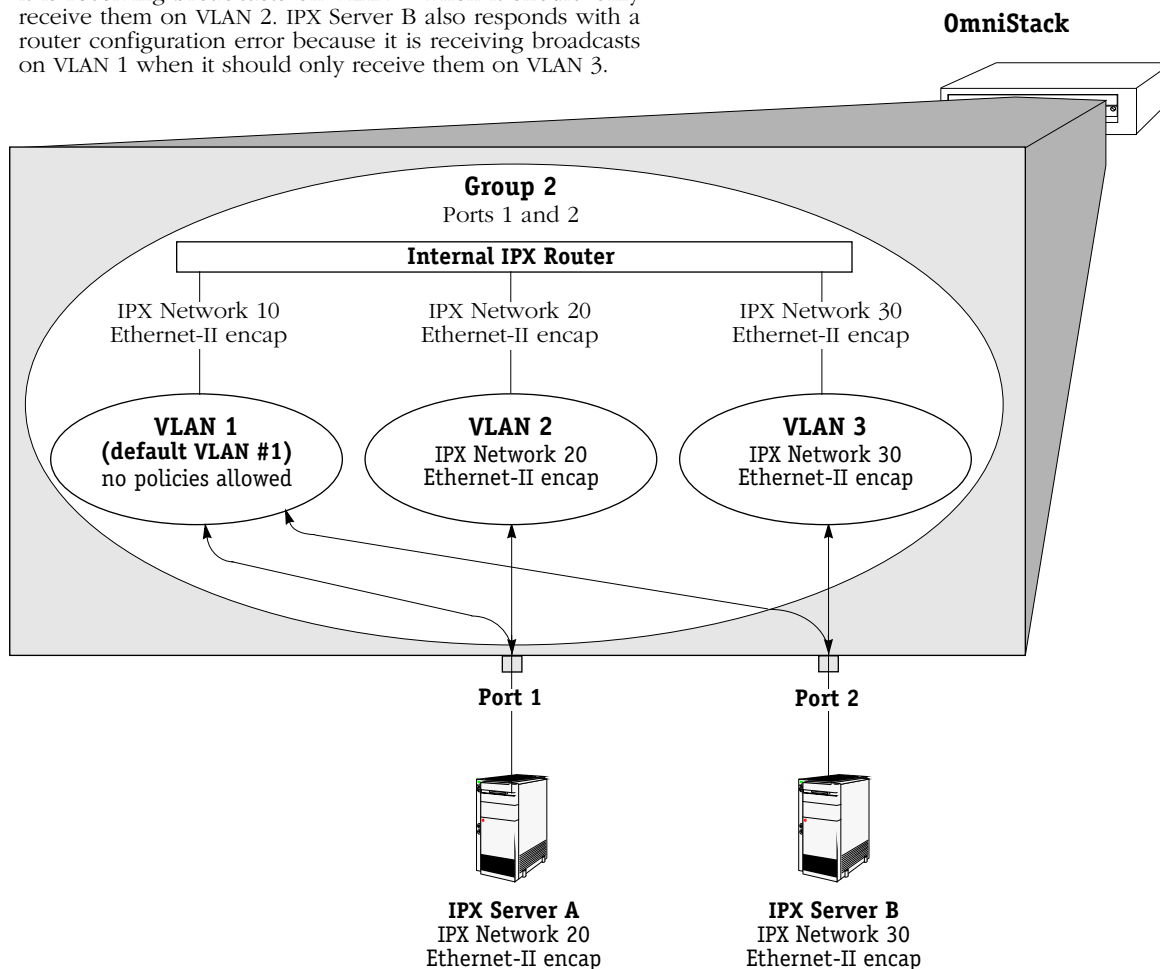
The Solution

The solution for this problem is to disable IPX routing on default VLAN #1. Because of this, when your network includes IPX servers that are members of IPX network address VLANs and IPX routing is enabled, you should configure your network such that disabling IPX routing on default VLAN #1 is not a problem.

Important Note

If you enable routing for a Group, you are actually enabling routing for that Group's default VLAN #1. For this reason, do not enable routing for any Group in which an IPX server is a member of an IPX network address VLAN.

When the internal IPX router sends out broadcasts on VLAN 1, they are flooded out all ports in the Group because, by default, all ports in the Group are members of VLAN 1. IPX Server A responds with a router configuration error because it is receiving broadcasts on VLAN 1 when it should only receive them on VLAN 2. IPX Server B also responds with a router configuration error because it is receiving broadcasts on VLAN 1 when it should only receive them on VLAN 3.



Application Example 5

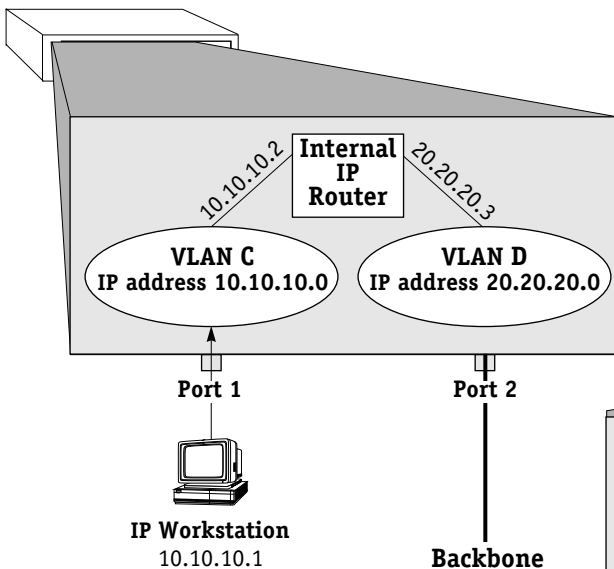
Traversing a Backbone

Application Example 5 illustrates why port-based policies may be required to establish communications in some network situations, such as traversing a backbone. This necessity arises because, as explained in *How Routing Works Generally* on page 22-8, AutoTracker does not activate a VLAN – or its internal router interface – 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.

The figure below illustrates the problem that can occur. The network below contains two OmniStacks in which three IP network address VLANs exist: VLAN C (IP address 10.10.10.0), VLAN D (IP address 20.20.20.0), and VLAN E (IP address 30.30.30.0). VLAN D spans both OmniStacks, but has no assigned devices. Routing is enabled for all three VLANs. A backbone connects port 2 on OmniStack 1 to port 1 on OmniStack 2.

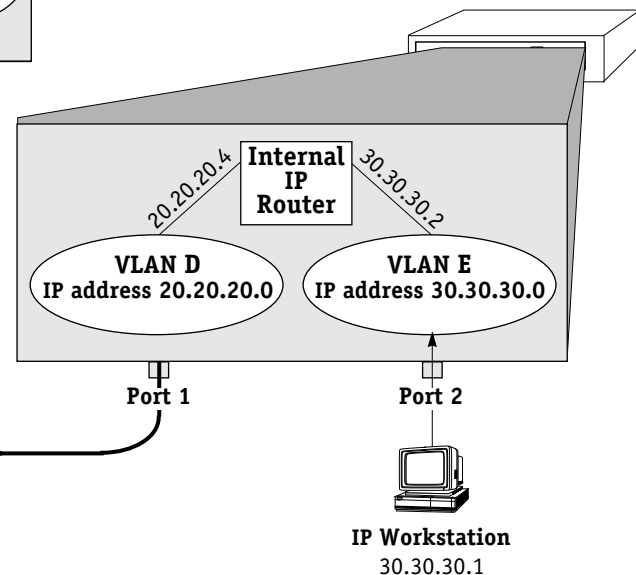
When IP workstation 10.10.10.1 transmits a frame VLAN C and its internal router activate. When IP workstation 30.30.30.1 transmits a frame VLAN E and its internal router activate. All subsequent traffic on VLAN C is transmitted to IP workstation 10.10.10.1 and all subsequent traffic on VLAN E is transmitted to IP workstation 30.30.30.1. VLAN D cannot activate because there are no devices that match its network address policy and it has no ports assigned. Because VLAN D is not active, Switches 1 and 2 cannot exchange routing information. Switch 1 will not be aware of network 30 and Switch 2 will not be aware of network 10.

OmniStack 1



With this configuration, VLAN D can never become active because it has neither assigned ports nor attached devices. Thus, Switches 1 and 2 cannot share routing information over the backbone.

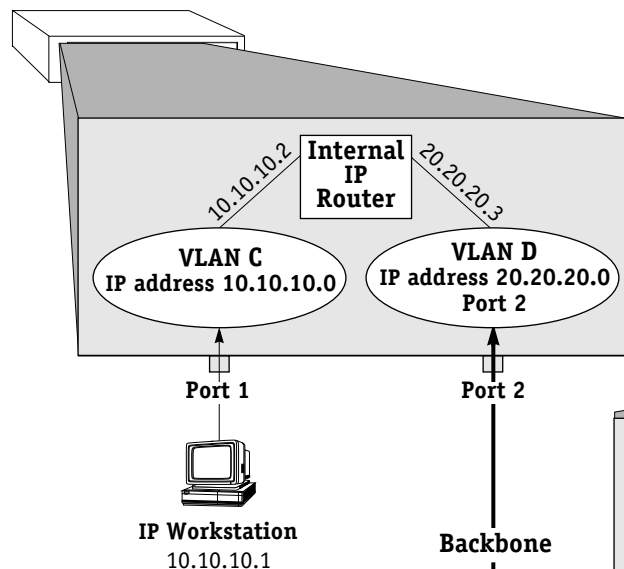
OmniStack 2



The Solution

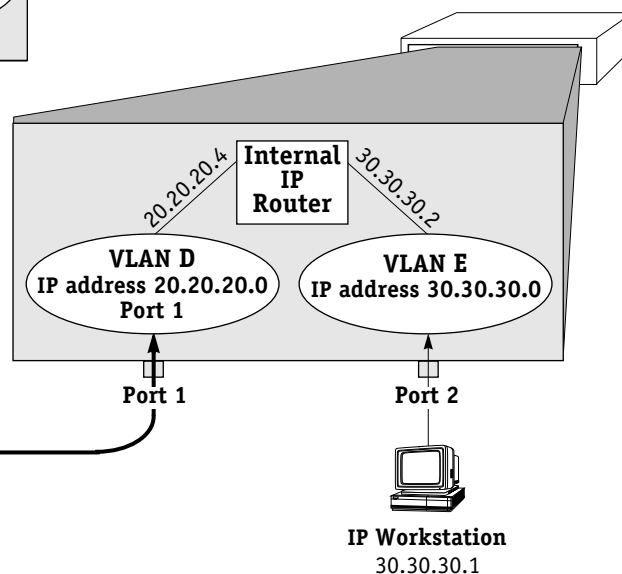
The recommended solution is to add a port policy to VLAN D, as is shown in the figure below. A port policy can be defined in addition to any other policies defined for a VLAN. If VLAN D has a port policy that includes port 2 on Switch 1 and port 1 on Switch 2 – the ports to which the backbone is connected – VLAN D and its internal router will activate immediately in both Switch 1 and Switch 2. Traffic (i.e., routing information) can then flow between Switch 1 and Switch 2 over the backbone. Switch 1 will be aware of network 30 and Switch 2 will be aware of network 10.

OmniStack 1



Adding a port policy to VLAN D that includes the ports to which the backbone is connected solves the problem. VLAN D now activates immediately – since it has ports assigned – and traffic can flow between the two switches.

OmniStack 2



Please Take Note

Refer to Chapter 17, “Configuring Group and VLAN Policies,” for information on original and current port policy functionality.

