

17 Managing Groups and Ports

In a traditional hub-based network, a broadcast domain is confined to a single network interface, such as Ethernet, or even a specific physical location, such as a department or building floor. In a switch-based network, such as one comprised on OmniStacks, a broadcast domain—or *Group*—can span multiple physical switches and can include ports using multiple network interfaces. For example, a single OmniStack Group could span three different switches located in different buildings and include Ethernet, ATM and WAN physical ports.

An unconfigured OmniStack contains one Group, or broadcast domain. It also contains one default Virtual Network, or VLAN, referred to as “default VLAN #1”. The default Group, Group #1, and its default VLAN contain all physical ports in the switch. When a switching module is added to the switch all of these additional physical ports are also assigned to Group #1, VLAN #1.

You can create Groups in addition to this default Group. When you add a new Group, you give it a name and number, optionally configure a virtual router port for its default VLAN, and then add switch ports to it. The switch ports you add to a new Group are moved from the default Group #1 to this new Group. (For more information on how ports are assigned to Groups, see *How Ports Are Assigned to Groups* on page 17-2.)

Up to 500 Groups can be configured on each OmniStack, but only 128 of those Groups can be active at one time. An entire OmniStack/OmniSwitch network can contain up to 65,535 Groups. Each Group is treated as a separate entity.

There are three main types of Groups:

1. **Mobile Groups.** These groups allow ports to be dynamically assigned to the Group based on AutoTracker polices. In contrast to non-mobile Groups, AutoTracker rules are assigned directly to a mobile Group. No AutoTracker VLANs are contained within a mobile Group. (However, mobile groups do contain a default VLAN 1 to which AutoTracker policies are assigned; policies assigned to this default VLAN apply to the entire mobile group.) Any AutoTracker policy may be used as criteria for membership in a mobile Group. Mobile groups are described in more detail in *Mobile Groups* on page 17-5.
2. **Mobile Groups based on authentication.** Authenticated Groups are a special form of mobile Group. These Groups include devices that are dynamically assigned based on an authentication criteria. Typically the user will have to log in with a valid password before being included in an authenticated mobile Group. Group membership is based on users proving their identity rather than the physical location of user devices. Authenticated Groups are described in more detail in the *Switch Network Services User Manual*.
3. **Non-mobile Groups.** These Groups contain statically assigned ports and may contain AutoTracker or Multicast VLANs. These VLANs within a non-mobile Group use AutoTracker policies to filter traffic. AutoTracker rules are not assigned to non-mobile Groups, they are assigned to the VLANs within the Group. Non-mobile groups are described in more detail in *Non-Mobile Groups and AutoTracker VLANs* on page 17-17.

All three types of Groups may co-exist on the same switch. However, a switch port cannot belong to a non-mobile group and a mobile group.

How Ports Are Assigned to Groups

There are two methods for assigning physical OmniStack ports to a Group. One method is static and requires manual configuration by the network administrator; the other method is dynamic and requires only the configuration of AutoTracker rules for port assignment to occur. The two methods are described in this section.

Static Port Assignment

In the static method, the network administrator manually assigns a port to a Group through the **crgp**, **addvp**, or **addqgp** commands. The static method can be restrictive because it limits the mobility of users in a multi-Group network. Users can only move within their assigned Group. In addition, customized access for individual users is limited by this method. You can use the static method of port assignment with mobile and non-mobile groups. Static port assignment can be combined with dynamic port assignment for mobile groups, while static port assignment is the only method for assigning ports to non-mobile groups.

Dynamic Port Assignment (Group Mobility)

The dynamic method is available with the Group Mobility feature. Initially each port is part of the default Group #1 (only ports in the default Group and ports in mobile Groups are candidates for dynamic port assignment). Based on the nature of traffic and configured AutoTracker policies, ports are dynamically assigned to the appropriate Group.

For example, if a device attached to a port transmits traffic from the 140.0.0.0 subnet, AutoTracker will check to see if a policy exists for this IP address. If it does, then it will move the port from the default Group to the first Group using this policy. If this device detaches from the network the port will be re-assigned to a Group without intervention by the network administrator.

A port can belong to multiple mobile groups (up to 16) as long as devices attached to that port match policies of these mobile groups. However, an individual device, or MAC address, can only belong to one mobile group per protocol.

The dynamic method of port-to-Group assignment still requires the creation of Groups through the **crgp** command. The criteria for the dynamic assignment of ports to a Group are determined by AutoTracker policies that you can configure during the **crgp** procedure.

Only Ethernet ports can be dynamically assigned to Groups.

If more than one Group has the same type of rule, then ports matching that policy will be assigned to the first Group matching the policy. For example, if a device matched policies in both Groups 2 and 5, the port would be assigned to Group 2. To make the most out of Group Mobility it is best not to duplicate policies among Groups.

Configuring Dynamic Port Assignment

You can enable dynamic port assignment while creating a group through the **crgp** command. During the **crgp** procedure, you will be prompted

Enable Group Mobility on the Group ? [y/n] (n):

Answer **Yes** to this question to give this Group the capability of having ports and devices dynamically added to the Group. Port and devices will be dynamically assigned based on AutoTracker rules you define.

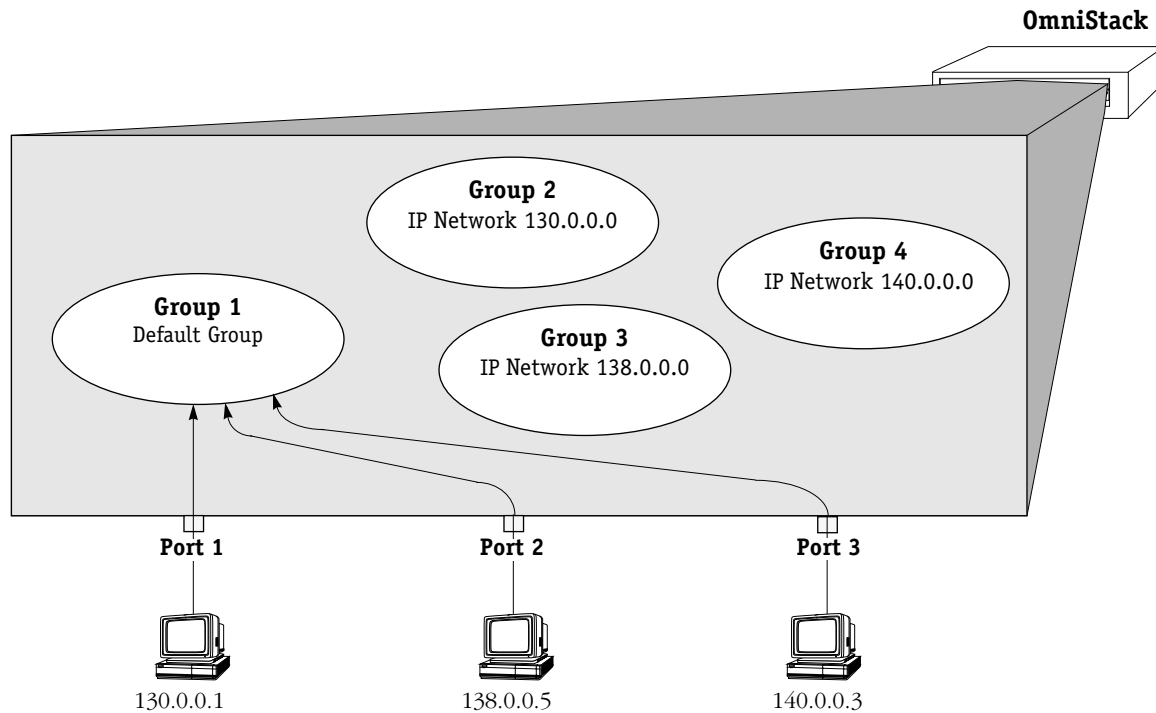
Service Ports and Group Mobility

Dynamic port assignment (ports carrying Ethernet traffic only) to Groups can also apply to LANE service ports configured for ATM access. These ports may be automatically added to the mobile group during the **crgp** procedure or through the **cats** command.

How Dynamic Port Assignment Works

Initially each port is assigned to the default Group. In this example, all three ports have workstations that belong to three different IP subnets (130.0.0.0, 138.0.0.0, and 140.0.0.0). All three ports start out in the default Group.

AutoTracker examines traffic coming from OmniStack ports. Three mobile groups are defined on the switch and each uses a different IP policy. Traffic that matches IP policies for a Group will trigger the movement of the port to the matching Group.



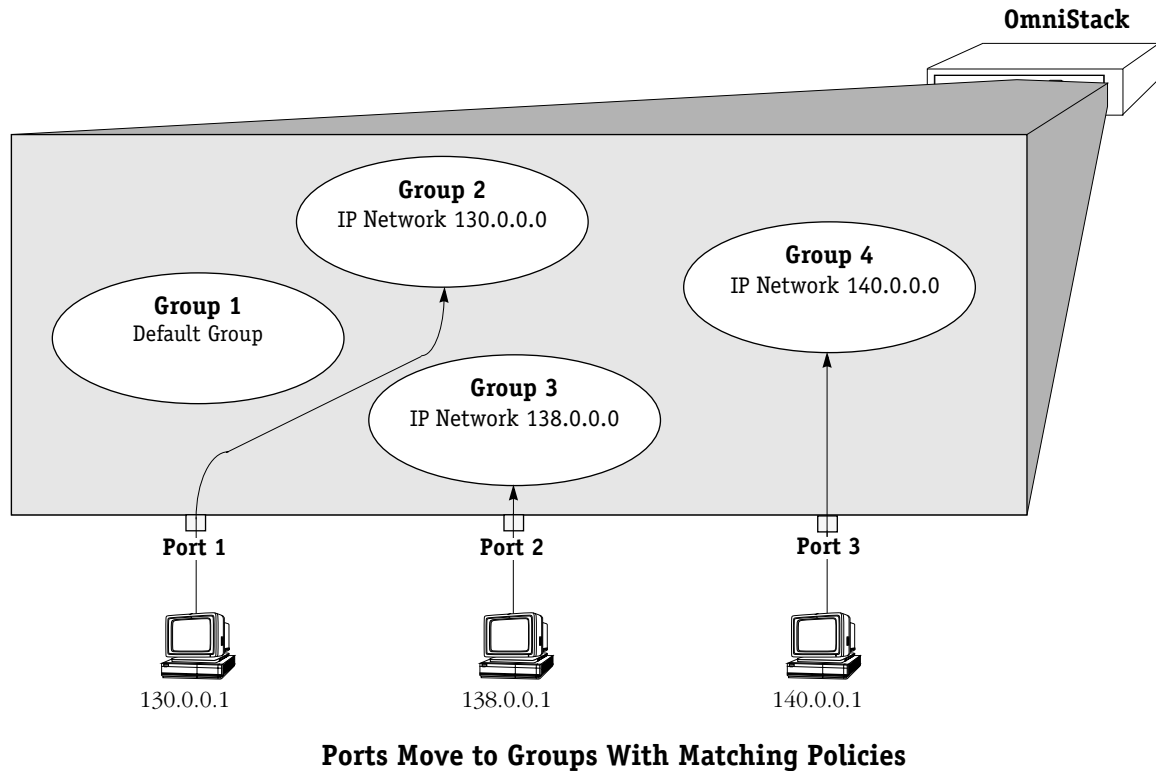
Initial Configuration: All Ports in Default Group

As soon as the workstations start transmitting traffic, AutoTracker checks the source subnet of the frames and looks for a match with any configured IP policies. If a match is found—and in this example all three ports can be matched with a corresponding Group—the port is moved to the matching Group.

Devices matching a policy trigger the assignment of a port to a mobile group. Therefore, the device is moved to the mobile group at the same time as the port to which it is attached. If more than one device comes in on a port, then that port can belong to more than one mobile group. Similarly, if a device transmits more than one protocol—such as IP and IPX—then the port to which it is attached can belong to more than one mobile group.

How Ports Are Assigned to Groups

As the illustration below shows, the three ports are each moved from the default Group to a Group with a policy that matches the subnet address of the workstation attached to the port. AutoTracker IP address policies have been set up in Groups 2, 3, and 4. The ports are moved to the Group with policies matching the subnet of the workstation.



Mobile Groups

Switch ports can be dynamically assigned to mobile groups through AutoTracker policies. Support for dynamic port assignment is one of the main differences between mobile groups and non-mobile groups. AutoTracker rules are assigned *directly* to a mobile group. In contrast, AutoTracker rules are assigned to the VLANs *within* a non-mobile group. No AutoTracker VLANs are contained within a mobile Group, and each mobile group constitutes a single spanning tree.

A switch port can belong to multiple mobile groups, whereas a switch port can belong to only one non-mobile group. However, a port can *not* belong to a mobile and a non-mobile group at the same time.

Ports can be assigned to mobile groups either statically or dynamically. A port is *statically* assigned to a mobile group when one of the following occurs:

- Port by default assigned to default group 1
- Port assigned to a group through **crgp** or **addvp** commands

Although switch ports can belong to multiple mobile groups, it is not possible to assign a port to two different groups using the **addvp** command. However, a switch port could be assigned to one mobile group via the **addvp** command and then gain membership to another mobile group by matching the policy criteria for that group.

A switch port is *dynamically* assigned to a mobile group after one of its attached devices matches an AutoTracker policy for that mobile group. An overview of how ports and devices are dynamically assigned to mobile Groups can be found in *How Ports Are Assigned to Groups* on page 17-2.

Dynamic LANE Services

Another feature of mobile groups is their support for dynamic LAN Emulation services. By specifying the ELAN name and the ATM port number of a LANE service, the LANE service will be brought up automatically as soon as Ethernet traffic is received from an Ethernet device. Dynamic LANE services are described in *Dynamic LANE Services* on page 17-14.

Authenticated Groups

Mobile groups provide the added flexibility of user-authentication policies. Using Authentication Management Console (AMC) software, you can configure mobile groups to use log-in procedures as a means of assigning group membership. Mobile groups that use authentication are a special group type called an Authenticated Group. Authenticated Groups are described in more detail in the *Switch Network Solutions User Manual*.

Configuring Mobile Groups

You configure mobile Groups through the **crgp** command. During the **crgp** procedure you will receive a prompt asking if you want to create a mobile Group

Enable Group Mobility on this Group ? [y/n] (n):

You must answer **Yes** to this prompt to set up a mobile group. After this question, you will be asked to configure virtual ports and AutoTracker policies for the Group. Documentation for the full **crgp** procedure can be found in *Creating a New Group* on page 17-20.

Turning Group Mobility On or Off

The **gmstat** command turns group mobility on or off for a Group that you specify. Essentially, you can change a non-mobile group into a mobile group and a mobile group back into a non-mobile group through **gmstat**. The group you specify must previously have been created through the **crgp** command.

Use the following syntax for the **gmstat** command:

```
gmstat <group number>
```

For example, if you wanted to change the group mobility status of group 2, you would enter:

```
gmstat 2
```

Mobile Group to Non-Mobile Group

If this group is already a mobile group, the following would display:

```
Group Mobility is ON for Group 2  
Change Group Mobility Status for Group 2 to OFF ? [y/n] (y):
```

If you wanted to change this mobile group back to a non-mobile group, you would press <enter> and the group would lose its mobile status. All AutoTracker policies you set up for the Group would no longer be valid.

If you decided not to turn off group mobility, enter **n** and the following prompt displays:

```
Group Mobility Status unchanged
```

Non-Mobile Group to Mobile Group

If this group is currently a non-mobile group, the following would display:

```
Group Mobility is OFF for Group 8  
Change Group Mobility Status for Group 8 to ON ? [y/n] (y):
```

If you wanted to turn on Group Mobility, you would press <enter> and would then be asked if you want to configure AutoTracker policies. If you answer yes, then the AutoTracker policies menu would display as follows:

```
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):
```

You define policies for a mobile Group. Non-mobile groups do not require policies (only the VLANs within them require policies). However, mobile Groups use policies to define membership. Instructions for specifying AutoTracker policies may be found in Chapter 20; please refer to that chapter for further instructions.

After you configure rules for the mobile group, you will be given the option to configure a dynamic LANE service. A description of this service is provided on page 17-14.

If you decided not to turn group mobility on, you would enter **n** at the group mobility prompt and the following message would display:

```
Group Mobility Status unchanged
```

Understanding Port Membership in Mobile Groups

Switch ports can belong to multiple mobile groups. A port becomes a member of a mobile group as long as one of its attached devices matches the policy criteria for that group. However, the movement of ports between groups and the status of port membership in groups can be affected by more than just whether or not devices match policy criteria.

Group mobility uses three variables that can affect a port's default group and whether or not a port ages out of a group. These variables are as follows: `def_group`, `move_from_def`, and `move_to_def`. The `def_group` and `move_to_def` variables can be configured through the `gmcfg` command, which is described on page 17-11. The `move_from_def` variable is enabled by default, but can be disabled by entering a statement in the `mp4.cmd` file. The effects of these three variables are described through diagrams on the following pages.

From the perspective of a device or switch port, there are three types of mobile group—default, primary, and secondary. Keep in mind that definitions of these three types are relative and can change for each port and device depending on the settings of the group mobility variables and traffic patterns of devices.

Default Group

The default group is the first group a port or device is assigned to by “default.” Typically, a port's default group will be Group 1. A port can also be statically assigned to its default group through the `crgp` or `addvp` commands. A port or device does not have to match a policy to gain membership into its default group.

The default group for a port or device is stored in memory; it can only be manually changed through the `addvp` or `crgp` commands. Depending on the settings of other group mobility variables a device or port can age out of other mobile groups but still remain a member of its default group.

Primary Group

The primary group is the group upon which Spanning Tree operations converge. The primary group is similar to the default group. In fact, the primary group for a port is initially the same as the default group. There are two main differences between a primary and a default group.

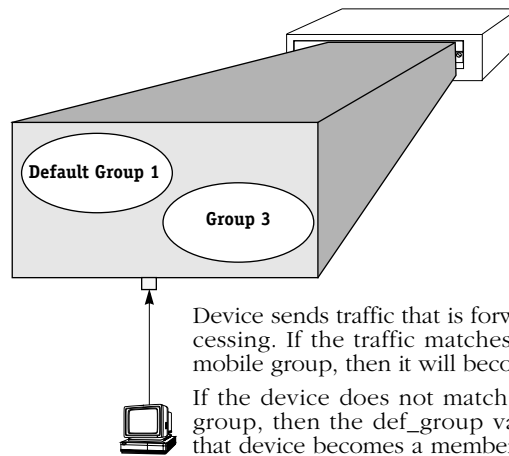
1. A primary group only contains devices that have matched one of its AutoTracker policies. In contrast, switch ports may end up in a default group without matching any policy.
2. It is possible for the primary group of a port or device to change through learning or aging. For example, if the `move_from_def` variable is enabled and a device matches the policies of a mobile group other than its default group, then this new mobile group becomes the primary group for the device and the port to which the device is attached (see diagram on page 17-9). In this case the default group and primary group will be different.

In addition a port can age out of its primary group if the `move_to_def` variable is enabled (see diagram on page 17-10). A port cannot age out of its default group.

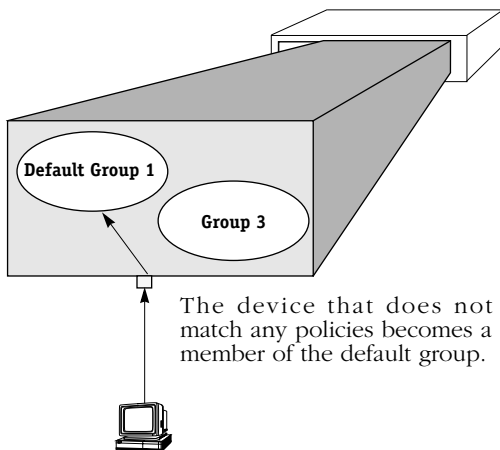
Secondary Group

Switch ports and devices may become members of multiple mobile groups. A switch port starts in its default group, which initially is also its primary group. The primary group may change if the `move_from_def` variable is enabled. Any subsequent mobile groups to which a port gains membership beyond the primary group are “secondary” mobile groups. A port can age out of these secondary groups if the `move_to_def` variable is enabled (see diagram on page 17-10).

How a Device Is Dropped from the Default Mobile Group (def_group)



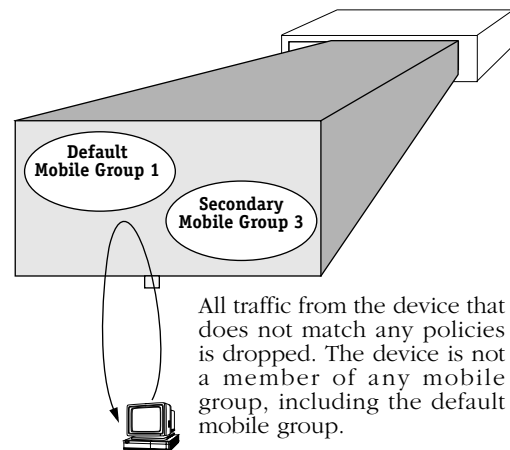
If `def_group` is enabled....



Why enable `def_group`?

- Ensure that all network devices will be a member of at least one mobile group.

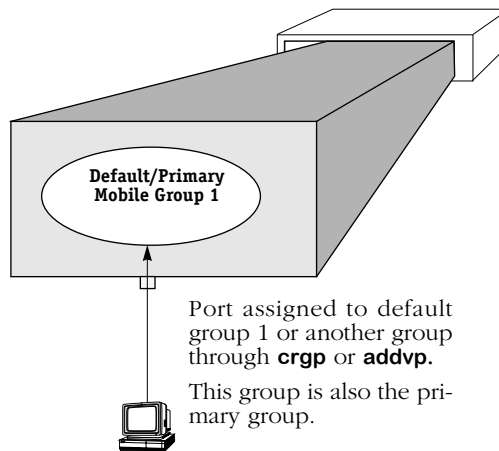
If `def_group` is disabled....



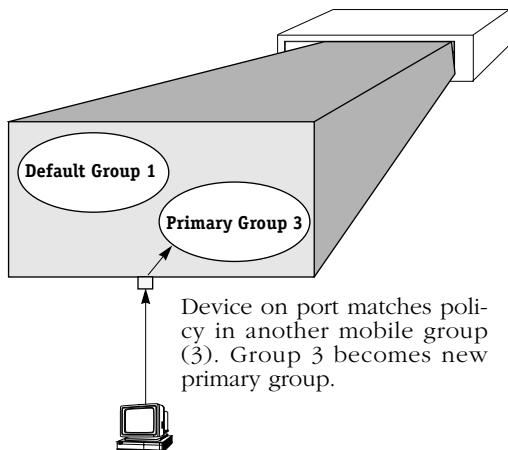
Why disable `move_from_def`?

- Reduces traffic to and from devices that do not satisfy any network policies.

How a Port's Primary Mobile Group Changes (move_from_def)



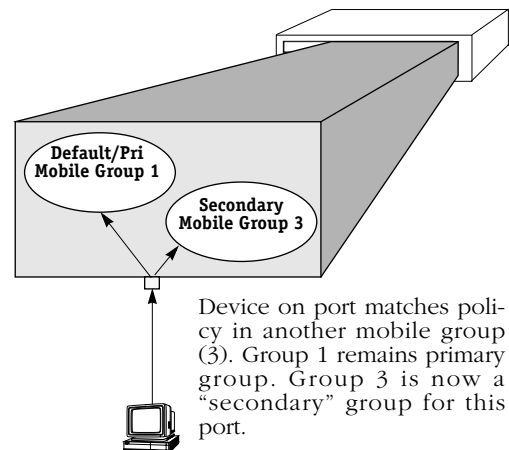
If move_from_def is enabled....



Helpful Hints:

- Reduces broadcasts to the default group.
- Best used when only one device is attached to each port.

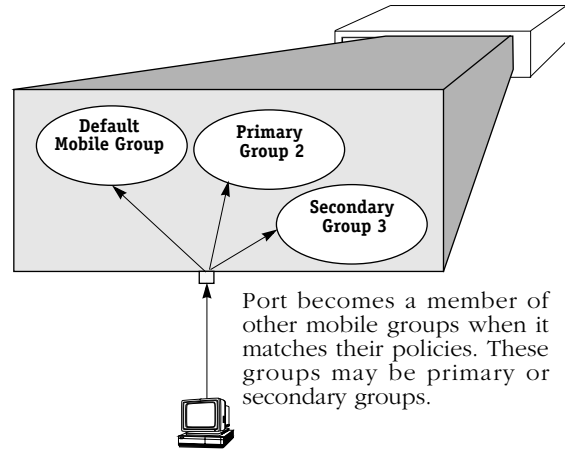
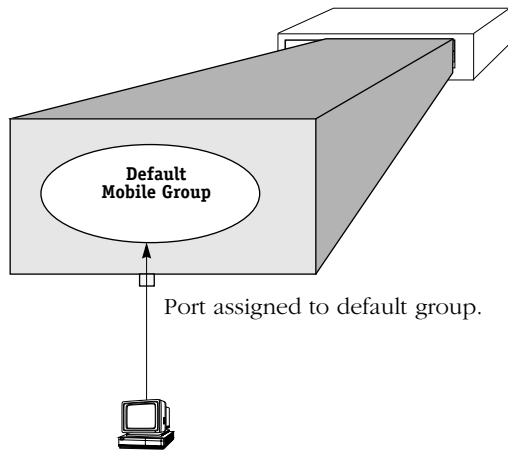
If move_from_def is disabled....



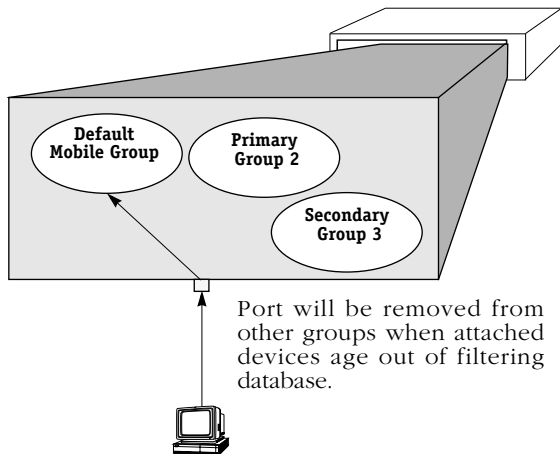
Why disable move_from_def?

- When multiple devices are attached to the switch port and the port must support multiple protocols, and thus multiple mobile groups.

How a Port Ages Out of a Mobile Group (move_to_def)



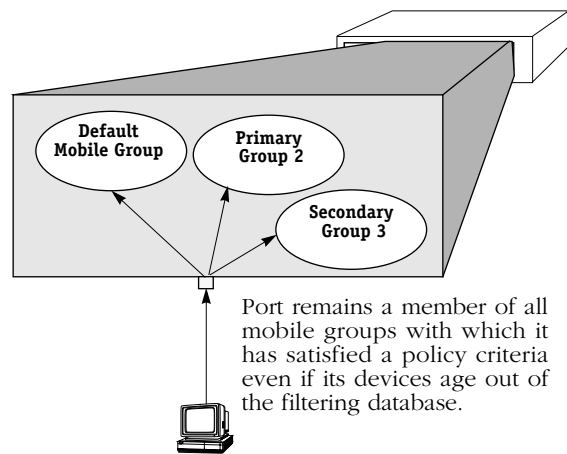
If move_to_def is enabled....



Why enable move_to_def?

- Security. Mobile groups only contain devices and ports that have recently matched policy criteria.

If move_to_def is disabled....



Why disable move_to_def?

- Switch ports retain group membership even when idle for some time. May be appropriate for silent devices, such as printers.

Configuring Switch-Wide Group Mobility Variables

There are several switch-wide group mobility variables that you can configure through the **gmcfg** command. These variables control the status of group mobility on all groups in a switch as well as the use of the default group. These variables are illustrated through diagrams on pages 17-8 to 17-10.

Follow these steps to use the **gmcfg** command:

1. Enter **gmcfg**. You do not need to specify a group number as this command applies to all mobile groups in this switch.
2. The following prompt displays:

Group Mobility is Enabled. Disable Group Mobility ? [yes/no] (no) :

This prompt controls the status of group mobility in this switch. If you disable group mobility here then mobile groups will not be supported in this switch even if they are configured through the **crqp** command.

Default Group 1. When group mobility is enabled, default group 1 in the switch will be treated as a mobile group and you will not be able to create AutoTracker VLANs within this group. When group mobility is disabled, default Group 1 in the switch will be treated as a non-mobile group in which AutoTracker VLANs could be created.

The default is to turn Group Mobility off. If you want to enable group mobility, then you need to indicate that choice at this prompt. The prompt will always show the current status of Group Mobility and then ask if you want to change that status. If you want to change the current status, then enter a **y** at this prompt and press <enter>. To keep the current status, simply press <enter>.

3. The following prompt displays:

move_to_def is set to Disabled. Set to Enable ? [yes/no] (no) :

The **move_to_def** variable determines what happens to a port once the devices on that port age out of the filtering database. By default this variable is Disabled, which means that a port will remain a member of a mobile group as long as its attached device satisfied the criteria for membership in that mobile group at one point. If devices on a port stop transmitting, the port will still retain all its mobile group memberships.

If the **move_to_def** variable is Enabled, then a port will lose its membership in a mobile group if its devices age out of the filtering database for that mobile group (i.e., they stop transmitting traffic that satisfies the criteria for membership in the mobile group). Once a port loses membership in all criteria-based mobile groups, it will return to its default group. The effect of this variable is illustrated on page 17-10.

By default, the **move_to_def** variable is Disabled. If you want to enable it (ports lose mobile group membership when they age out), then you need to indicate that choice at this prompt. The prompt will always show the current status of **move_to_def** and then ask if you want to change that status. If you want to change the current status, then enter a **y** at this prompt and press <enter>. To keep the current status, simply press <enter>.

4. The following prompt displays:

def_group is set to Enable. Set to Disable ? [yes/no] (no) :

The **def_group** variable determines what happens to devices that do not match any mobile group policies. If **def_group** is Enabled (the default), then devices that do not match any mobile group policies will be part of the default group for that port. If the **def_group** variable is Disabled, then devices that do not match any mobile group policies will be dropped from their default group and will not be part of any mobile group.

By default the `def_group` variable is Enabled. If you want to disable it (devices that do not meet criteria for mobile group membership will not be part of any mobile group), then you need to indicate that choice at this prompt. The prompt will always show the current status of `def_group` and then ask if you want to change that status. If you want to change the current status, then enter a **y** at this prompt and press **<enter>**. To keep the current status, simply press **<enter>**.

The `move_from_def` Variable

The `move_from_def` variable controls whether or not a port's primary group can differ from the port's default mobile group. This variable is enabled by default, but can be changed to disabled in the **mp4.cmd** file.

The original default group for a port is group 1 or the group to which the port is assigned through the **crgp** or **addvp** commands. The primary group at this point is the same as the default group. However, if the `move_from_def` variable is enabled, the primary group can change as soon as a device on the port matches the policy criteria for another mobile group.

For example, Port 5 may start out in Group 1, its default group. The primary group in this case will also be Group 1. If the `move_from_def` variable is enabled and Port 5 matches AutoTracker policies for mobile group 3, then the new primary group for Port 5 will be Group 3. All further Spanning Tree operations for the port will converge on group 3 rather than group 1. The effects of the `move_from_def` variable are further illustrated through diagrams on page 17-9.

If you disable the `move_from_def` variable, then the primary group for a port will always match the default group regardless of the number of other mobile groups to which it gains membership. To disable the `move_from_def` variable, enter the following statement in the **mp4.cmd** file

```
move_from_def=0
```

For this new setting to take place you need to reboot the switch.

Viewing Ports in a Mobile Group

The **vpl** command lists all the Groups in the switch currently configured as mobile Groups and the ports currently assigned to those Groups. Since ports are assigned to mobile groups dynamically, this display is helpful to find out which ports the switch already sees in each group. Ports will only display in this screen for secondary groups (i.e., not default or primary groups). Enter **vpl** and a screen similar to the following displays:

```
=====
Group ID      Physical Port      Virtual Port
=====
Group ID: 2    4/2 4/3 4/4 4/5    12 13 14 15
Group ID: 3    3/1 5/2            8 20
Group ID: 6    NULL Port List
Group ID: 8    4/1 5/1            11 19
=====
```

Group ID. The group number assigned to this mobile group during the **crgp** procedure.

Physical Port. The physical switch ports that have been dynamically assigned to this group because they matched an AutoTracker policy. (Primary groups do not display in this screen. For a display of port-to-primary group mappings, use the **vi** command) If this column reads **NULL Port List**, then no physical ports have been assigned to the group yet.

Virtual Port. The virtual ports that are part of this mobile group. For Ethernet switch ports, there is a one-to-one relationship between physical and virtual ports. For ATM ports, multiple virtual ports may be associated with one physical port.

Viewing a Port's Mobile Group Affiliations

The **vigl** command lists all the ports in the switch that have been assigned to mobile Groups. It is similar to the **vpl** command, but it lists ports first and then Groups. Since ports are assigned to mobile groups dynamically, this display is helpful to find out which ports the switch already sees in each group. Ports will only display in this screen for secondary groups (i.e., not default or primary groups). Enter **vigl** and a screen similar to the following displays:

```
=====
Virtual Port   Physical Port      Group ID
=====
12 13 14 15    4/2 4/3 4/4 4/5    Group ID: 2
8 20           3/1 5/2            Group ID: 3
NULL Port List                               Group ID: 6
11 19          Physical Port      Group ID
=====
```

Virtual Port. The virtual ports in this mobile group. For Ethernet switch ports, there is a one-to-one relationship between physical and virtual ports. For ATM ports, multiple virtual ports may be associated with one physical port.

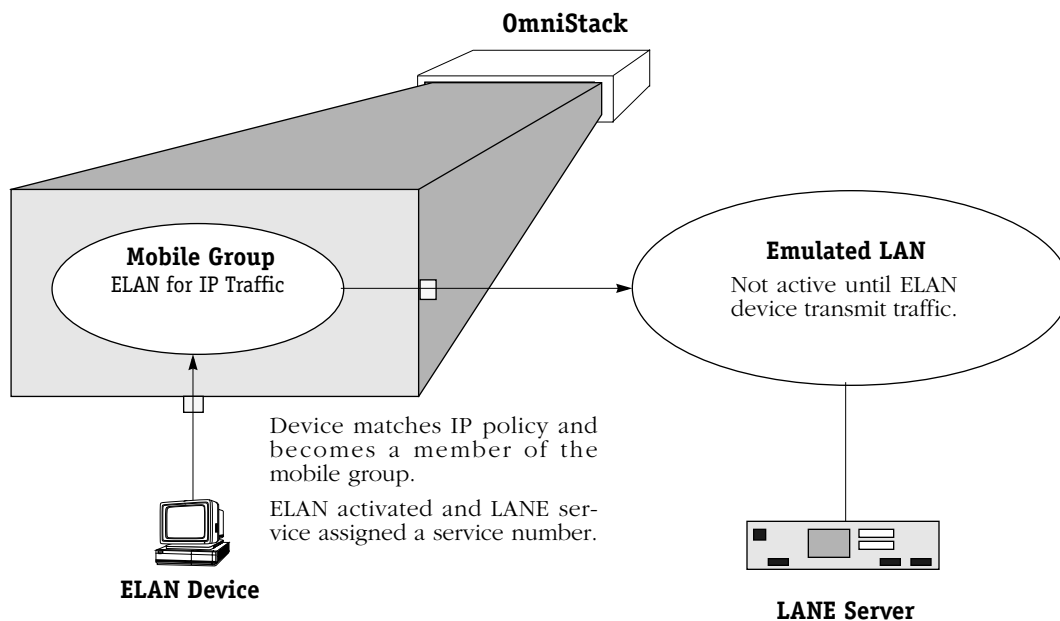
Physical Port. The physical switch ports that have been dynamically assigned to this secondary mobile group because they matched an AutoTracker policy. (Primary groups do not display in this screen. For a display of port-to-primary group mappings, use the **vi** command) If this column reads **NULL Port List**, then no physical ports have been assigned to the group yet.

Group ID. The group number assigned to this mobile group during the **crgp** procedure.

Dynamic LANE Services

Mobile groups allow you to create dynamic LAN Emulation (LANE) services, which can be a powerful tool in reducing configuration time and LANE broadcasts. By configuring a dynamic LANE service within a mobile group, you enter basic configuration information for the ELAN but the ELAN is not activated until one of the attached devices sends traffic.

Normally when you create a LANE service (through the **cas** command) all switch ports in the associated group receive all broadcasts associated with an ELAN. Even switch ports with no devices currently active in an ELAN will receive these broadcasts. However, when you create a dynamic LANE service the service is not activated until a switch port receives traffic from a member device.



How a Dynamic LANE Service is Activated

Dynamic LANE services function the same way as standard LANE services created through the **cas** command. The difference is that they are only activated when traffic is sent by an Ethernet device. In this sense, a dynamic service can be considered an on-the-fly version of the **cas** command. Configuration information on the ELAN is retained, but the service itself is dormant until an Ethernet device transmits data.

Dynamic LANE Services and Non-Dynamic LANE Services

Once an Ethernet device sends data, the LANE service will be activated by the switch and assigned a service number. The service, once activated, can be viewed through the **vas** command or modified through the **mas** command—just like any other ATM service. You can also delete a dynamic service with the **das** command; the only limitation to deleting a dynamic service this way is that the service only goes dormant and can still become re-activated by Ethernet traffic. A more foolproof way to delete a dynamic LANE service is through the **dats** command, which is described on page 17-15.

The ELAN name configured for the dynamic service cannot be the same as an ELAN configured through the **cas** command. If you create a service through **cas** that uses the same name as a dormant dynamic service, then the dynamic service will not become operational.

Creating Auto-Activated LANE Services

You can add a LAN Emulation service to a mobile group. The LANE service will automatically be created in the mobile group as soon as Ethernet traffic is received. You can add a dynamic service to a mobile Group during the **crgp** procedure (see *Step 5. Configure Auto-Activated LANE Ports (Mobile Groups Only)* on page 17-33) or by using the **cats** command. The LANE service is dormant until traffic occurs, at which time the service is activated and ports are assigned to the service.

Follow these steps to use the **cats** command:

1. Enter **cats**. The following prompt displays

Enter the mobile Group Id:

You can also enter the **cats** command followed by the group number to bypass this prompt.

2. Enter the group number for the mobile group to which you want to assign this LANE service port. The following prompt displays:

Enter the primary slot/interface for this service:

3. Descriptions of this prompt and the remaining **cats** prompts can be found in the **crgp** procedure. Please continue with the section, *Step 5. Configure Auto-Activated LANE Ports (Mobile Groups Only)* on page 17-33.

Deleting an Auto-Activated Service

You can delete any dynamic LANE services that you assigned to a mobile group. These services are assigned through the **crgp** or **cats** commands. You use the **dats** command to delete these services. Note that the **dats** command deletes only auto-activated services; it does not delete services created via the **cas** command.

Note

You could also delete an auto-activated service through the **das** command. However, **das** only renders an auto-activated service dormant; the service could re-activate through activity on a member port.

You will need to know the mobile group number and the service number of the auto-activated service before you can delete it. You can obtain the service number through the **vats** command, which is described in *Viewing Auto-Activated Services* on page 17-16.

The **dats** command uses the following syntax:

dats <group number> <slot>/<port> <service index>

For example, if you wanted to delete the LANE service with an index number of 3 on port 4/2 in mobile group 8, you would enter:

dats 8 4/2 3

The service port is deleted once you enter the command and press **<enter>**. You can obtain the service index number for a dynamic LANE service through the **vats** command.

Viewing Auto-Activated Services

Using the **vats** command, you can obtain a listing of any LANE services that have been assigned to a mobile group. These ports are assigned through the **crgp** or **cats** commands. Note that the **vats** command displays only auto-activated services; it does not show services created via the **cas** command.

When you enter **vats**, a screen similar to the following displays:

Group ID	P. Port S. Port A. Port	Srvc Num	Srvc Index	Admin	Oper	Srvc Type	Service Name
3	3/1	2	1	Enabled	Inactive	LANE	elan1
8	4/2	3	1	Enabled	Active	LANE	elan2
	4/1						
	4/2						

Group ID. The mobile group number to which this service is assigned.

P. Port, S. Port, A. Port. This column lists the current primary physical port (**P. Port**), secondary port (**S. Port**), and the current active port (**A. Port**) for this LANE service.

Srvc Num. The internal service number assigned to this service. You will need this number if you want to modify or delete a service through the **mas** or **das** commands, respectively.

Srvc Index. The internal value used by dynamic LANE software to identify this service. The service index differs from the service number in that it is used exclusively by dynamic LANE software. You will need the service index to delete a dynamic service through the **datas** command.

Admin. The current administrative status of the service. When configuring auto-activated services through **crgp** or **cats**, you have the choice to enable or disable the service. If **Enabled**, the service may become operationally active if traffic uses this service and the connection is good. If **Disabled**, the service cannot become active until you enable it.

Oper. the current operational status of the service. If **Active**, the service is currently in use. If you have set up primary and secondary services, only one will be Active. A port cannot become operationally active until it has been administratively enabled.

Srvc Type. The type of ATM service. Currently only LANE services are supported as auto-activated services.

Service Name. The ELAN to which this service belongs.

Non-Mobile Groups and AutoTracker VLANs

Non-mobile Groups are comprised of *physical* entities—switch ports. Groups can span multiple switches, but they are still made up of physical ports that you can see and touch. But just as physically-based broadcast domains are limited, entirely port-based Groups can also be limiting. In a large, flat, switched network, broadcast traffic can overload the network. There needs to be a method for subdividing traffic even further. That's where virtual networks, or VLANs, come into play.

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. VLANs are described in more detail in Chapter 20, “Managing AutoTracker VLANs” and Chapter 21, “Multicast VLANs.”

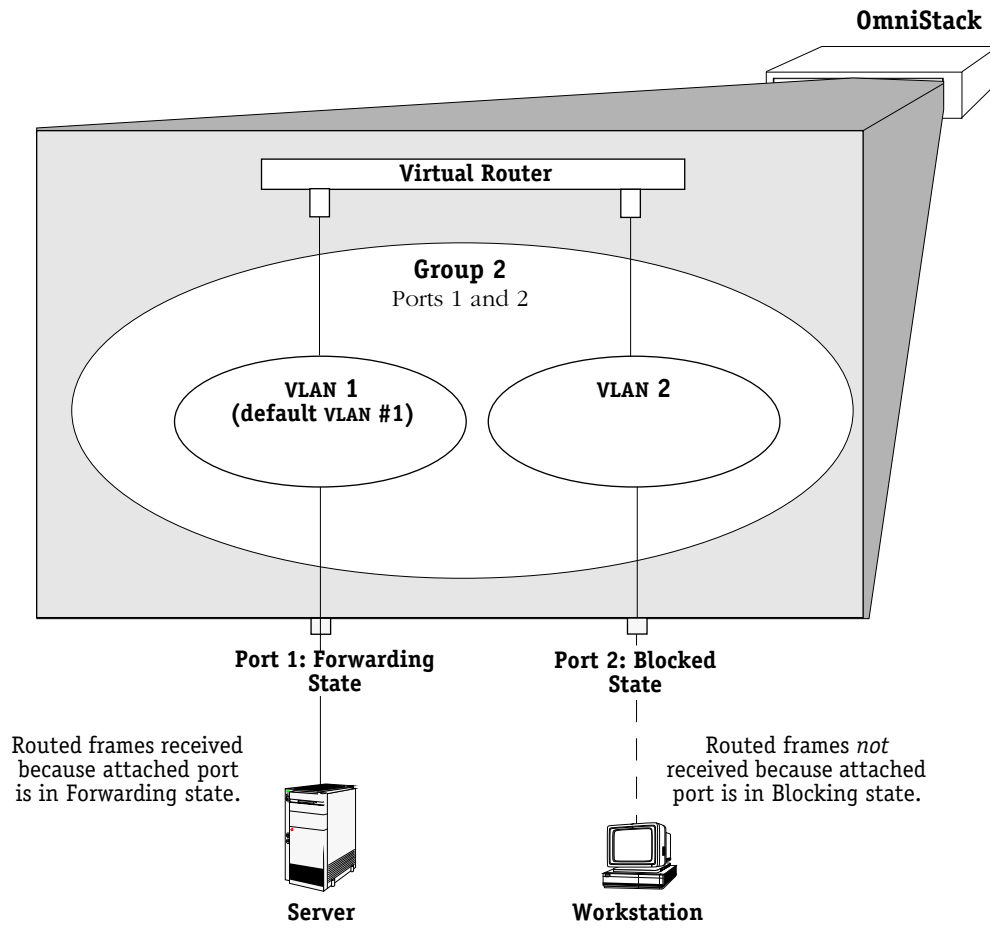
Routing in a Non-Mobile Group

Communication within a Group containing only the default VLAN is switched; the ports are in the same broadcast domain and do not require routing to communicate. Communication between VLANs in the same Group or to VLANs in other Groups requires routing. That's why all VLANs—including the default VLAN within each Group—may contain their own virtual router port. A virtual router port for each VLAN can be configured to support IP and/or IPX routing. If you do not configure a virtual router port for a VLAN, the devices in that VLAN will not be able to communicate with devices in other VLANs unless there is an external router between the VLANs.

Each OmniStack supports up to 32 virtual router ports. A single router port, using one MAC address, can support IP routing, IPX routing, or both types of routing. When you enable a router port for a default VLAN, you are actually creating a static route to that VLAN. Routing is covered in more detail in Chapters 23 and 25.

Spanning Tree and Non-Mobile Groups

Each Group uses one Spanning Tree for bridging. The OmniStack supports both 802.1d and IBM Spanning Tree protocols. The Spanning Tree state for the port is Forwarding. Ports that are in Blocked state, or in another non-Forwarding state, will not receive frames from the router port. The figure below illustrates this concept.



Spanning Tree State and Routed Frames

Group and Port Software Commands

Group and Virtual Port commands are part of the VLAN menu within the User Interface. Entering **vlan** at any prompt displays the following menu:

Command	VLAN Management Menu
gp	View the list of Groups currently defined
crgp	Create a Group
modvl	Modify a VLANs configuration/availability
rmgp	Remove a Group
addqgp	Add 802.1q group/s to a port
delqgp	Delete 802.1q group/s from a port
viqgp	Display 802.1q groups on port/s
via	View ports assigned to the selected Group
vi	View info on a specific virtual port
vs	View statistics on a virtual port attachment
ve	View errors on a virtual port attachment
addvp	Add ports to a GROUP
modvp	Modify existing VPORT configuration information
rmvp	Remove ports from a Group
pmapcr	Create a Port Map
pmapdel	Delete a Port Map
pmapmod	Modify a Port Map
pmapv	View Port Mapping Configuration
br	Enter the Bridge Configuration/Parameter sub-menu
prty_mod	Modify the priority of a group
prty_disp	Display the priority of a group
at	Enter the AutoTracker sub-menu

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

The VLAN menu commands are divided into four sets of commands. The first set, at the top of the menu beginning with **gp**, contains commands that create, modify, delete, and view Groups. The second set of commands, beginning with **addqgp**, allow you to manage 802.1q groups and display various configuration, status, and statistical information on virtual ports in the switch. (The 802.1q commands are described in Chapter 14.) The third set, beginning with **addvp**, contains commands for adding, modifying, and deleting virtual ports. All of these commands are described in this chapter.

The final set of commands at the bottom of the menu, **br** and **at**, are actually entry points to the Bridging and AutoTracker submenus, respectively. Commands for the Bridge Management (**br**) sub-menu are documented in Chapter 15, “Configuring Bridging Parameters.” Commands for the AutoTracker (**at**) sub-menu are documented in this chapter and in Chapter 20, “Managing AutoTracker VLANs” and Chapter 21, “Multicast VLANs.” Some commands in the **at** sub-menu apply to mobile groups and authenticated groups; those commands are described in this chapter.

The **pmapcr**, **pmapdel**, **pmapmod**, and **pmapv** commands allow you to create port mapping configurations. The port mapping feature is documented in *Port Mapping* on page 17-67. The **prty_mod** and **prty_disp** commands allow you to modify and view the priority of a selected group. These commands are detailed in *Priority VLANs* on page 17-74

Creating a New Group

There are several steps involved in creating a new Group. Note that some steps apply only to mobile groups. These steps are as follows:

1. Enter Basic Group Information, such as the Group number and type. This section starts on page 17-21.
2. Configure the Virtual Router Port (Optional). This section starts on page 17-22.
3. Enable/disable Group Mobility and User Authentication. This section starts on page 17-27.
4. Configure Virtual Ports. This section starts on page 17-28.
5. Configure Auto-Activated LANE service ports (for mobile groups only). This section starts on page 17-33.
6. Configure AutoTracker policies (for mobile groups only). This section starts on page 17-34.

WAN Routing and ATM Classical IP (CIP) Groups follow a slightly different procedure for their creation. You will receive prompts during the procedure asking whether you want to create one of these special Groups.

Step 1. Entering Basic Group Information

- a. Type **crgp** at any prompt. You can also enter the group number after the **crgp** command or allow the system to assign a number for you.

- b. The following prompt displays:

GROUP Number (5):

By default the Group number you entered or the next available Group number is displayed in parentheses. Enter the Group number or accept the number shown in parentheses. Each Group must have a unique number, which may range from 2 to 65,534. (Group 1 is the default switch Group. It does not need to be created and it cannot be deleted.) Press **<Enter>** after entering the Group number.

- c. The following prompt displays:

Description (no quotes) :

Enter a descriptive name for the new Group. Group names can consist of up to 30 alphanumeric characters. Press **<Enter>** after entering the Group name.

- d. The following prompt displays:

Enable WAN Routing? (n):

If you want to perform WAN Routing through this Group you must enter a **y** at this prompt. If you do not need to support WAN Routing, then answer **n** at this prompt and continue with Step e.

Note

You do not need to create a special WAN Routing Group to bridge or trunk traffic over a WAN connection. If you are just Bridging or Trunking on WAN, answer **n** to this prompt and continue with Step e.

A WAN Routing Group is different from other Groups; it must contain only WAN ports. In addition, the virtual router and virtual ports are configured differently. Please skip ahead to *Creating a WAN Routing Group* on page 17-35 to continue setting up this WAN Routing Group.

- e. The following prompt displays:

Enable ATM CIP? (n):

ATM Classical IP (CIP) must be enabled if there are devices on the ATM network that support only Classical IP or require Classical IP to communicate over the network. Answer **n** at this prompt if this Group will not support ATM CIP and skip ahead to *Step 2. Configuring the Virtual Router Port (Optional)* on page 17-22.

Answer **y** at this prompt if this Group must support ATM CIP. If you answer **y**, then this Group will support only CIP and you will need to configure it further through the Services menu. Skip ahead to *Creating an ATM CIP Group* on page 17-38 for further instructions on setting up this ATM CIP Group.

Step 2. Configuring the Virtual Router Port (Optional)

You can now optionally configure the virtual router port that the default VLAN in this Group will use to communicate with other VLANs. When you define a virtual router, a virtual router port for the default VLAN in the Group is created. If you do not define a virtual router, no virtual router port is created and the default VLAN in the new Group will be “firewalled,” unable to communicate with other VLANs.

◆ Important Note ◆

Use caution when setting up routing on the default VLAN for a Group. In some configurations enabling routing on the default VLAN may not be necessary or desirable. You can always enable routing on other, non-default VLANs, within this Group. Refer to *AutoTracker Application Example 4* in Chapter 22 for further information.

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

- a. After answering **n** to the **Enable ATM CIP?** prompt, 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 the default VLAN in this Group will not be able to route IP data. If you don't want to set up an IP router, enter **n**, press **<Enter>** and skip to Step k.

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.

- b. 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 of the default VLAN within this Group. After you enter the address, press **<Enter>**.

- c. The following prompt displays:

IP Subnet Mask (0xfffff00):

The default IP subnet mask (in parentheses) is automatically derived from the default VLAN 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>**.

- d. The following prompt displays:

IP Broadcast Address (198.200.10.255):

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

- e. 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>**.

- f. The following prompt displays:

Configure as Loopback? (n) :

Indicate whether you want to configure the group in loopback mode. Typically a group should not be configured for loopback unless you are running Advanced Routing software such as GateD.

- g. The following prompt displays:

Disable routing? (n) :

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

- h. The following prompt displays:

Enable NHRP? (n) :

Indicate whether you want to enable NHRP.

- i. 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 the default VLAN in this Group 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 the default VLAN in this Group 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. The default VLAN in this Group 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.

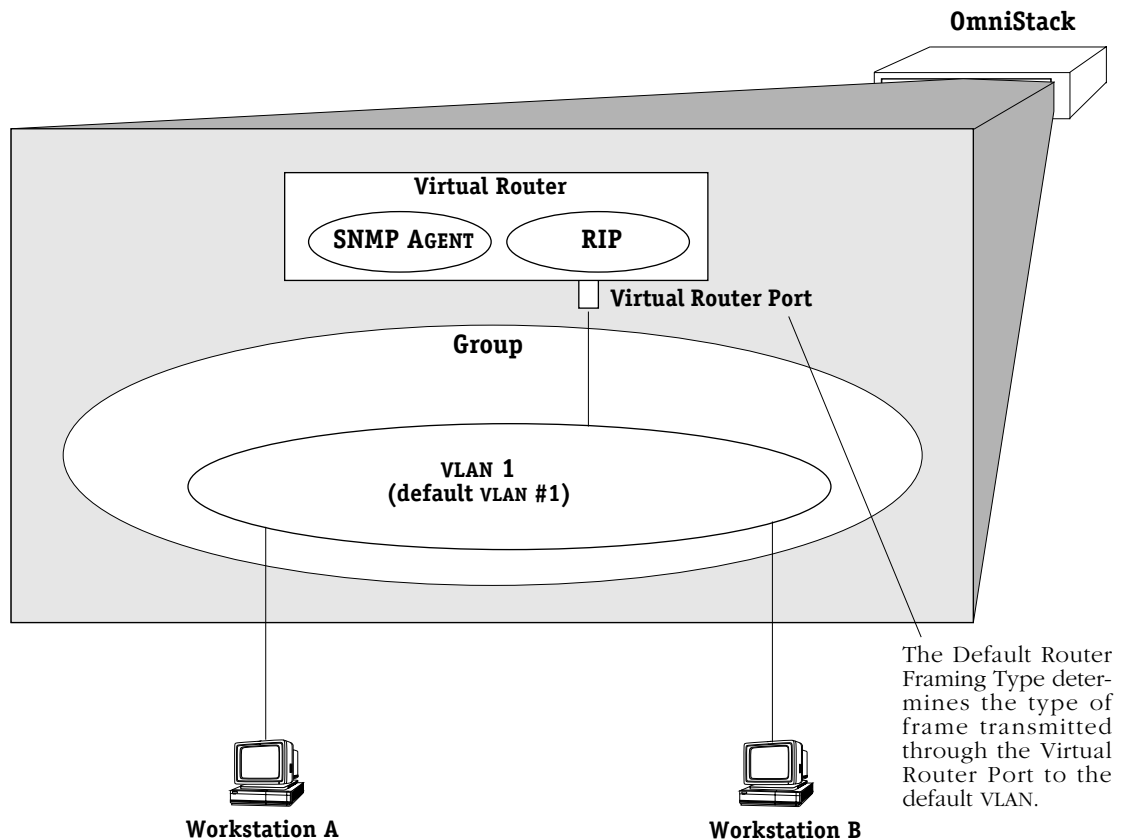
Active. RIP is active and both sends and receives RIP updates. The default VLAN in this Group will receive routing information from other VLANs and will be included in the routing tables of other VLANs. Enter **a** and press **<Enter>** to select Active mode.

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

- j. 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 the default VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in the default VLAN. If the default VLAN contains devices using encapsulation types other than those defined here, the switching modules 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.



Default Framing Type and the Virtual Router Port

- k. 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 the default VLAN in this Group will not be able to route IPX data. You can set up a virtual router port to route both IP and IPX traffic.

If you don't want to set up an IPX router for the default VLAN in this Group, enter **n**, press **<Enter>**, and skip ahead to step **p**. You can always set up IPX routing for other VLANs within this Group.

- l. 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.

- m. 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>**.

- n. 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 information protocol) and SAP (service access protocol), to be configured for the default VLAN in this Group. 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 default 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 default 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 default VLAN to which it is attached does not participate in either RIP nor SAP updates. Enter an **i** and press **<Enter>** to select RIP and SAP inactive.

- o. 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 default VLAN to the outbound ports. Set the framing type to the encapsulation type that is most prevalent in the default VLAN. If the default VLAN contains devices using encapsulation types other than those defined here, the switching modules must translate those frames, which slows throughput. See the figure, *Default Framing Type and the Virtual Router Port* on page 17-24 for an illustration of the Default Framing Type and its relation to Virtual Router Port communications.

Note

If you are using multiple encapsulations (for example, VLAN 2:1 is 802.3 and VLAN 3:1 is Ethernet II) then the **hrexnative** command must be set to zero in the **mp4.cmd** file. See Chapter 5, “Managing Files,” for information about modifying files.

- p. If you chose a Source Routing frame format in the last step (options 5, 7, 9, or b), 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.

- q. The following prompt displays:

Enter a priority level (0...7)(0):

Prioritizing VLANs allows you to set a value for traffic based on the destination VLAN of packets. Traffic with the higher priority destination will be delivered first. VLAN priority can be set from 0 to 7, with 7 being the level with the most priority.

Modifying and displaying a group's priority is described in *Priority VLANs* on page 17-74.

You have now completed the configuration of the virtual router port for this group. At this point, you will be asked whether you want to enable group mobility. The following prompt will display:

Enable Group Mobility on the Group ? [y/n] (n):

Mobile groups are discussed in detail in *Mobile Groups* on page 17-5. If you want to enable group mobility answer **Y** to this prompt, press **<enter>**, and go on to *Step 3. Set Up Group Mobility and User Authentication* on page 17-27.

If you do not want to configure group mobility answer **N** at the prompt, press **<enter>**, and go on to *Step 4. Configuring Virtual Ports* on page 17-28 for further instructions.

Step 3. Set Up Group Mobility and User Authentication

A mobile group offers more flexibility than a non-mobile group. With a mobile group, ports are assigned dynamically to the group based on AutoTracker policies that you configure. In a non-mobile group, ports are statically defined and AutoTracker policies are assigned to individual VLANs within the Group. In most cases, you will want to set up a mobile group. The following steps show you how.

- a. After configuring the virtual router port, you will receive the following prompt:

Enable Group Mobility on the Group ? [y/n] (n):

To create a mobile group, enter a **Y** as this prompt, press **<enter>**, and continue with step b. If you want to configure a non-mobile Group, enter **N**, press **<enter>**, and you will see the following prompt:

This Group will not participate in Group Mobility

If you are *not* creating a mobile group, go on to *Step 4. Configuring Virtual Ports* on page 17-28.

- b. The following prompt displays:

Enable User Authentication on the Group ? [y/n] (n):

An authenticated group is a special type of mobile group. It uses an authentication process as its criteria for group membership. Typically, users will be prompted for an id and password before gaining membership to an authenticated group. Authenticated groups require additional Windows NT server software. More detailed information on these groups can be found in the *Switch Network Solutions User Manual*. If you are not sure whether this is an authenticated group, simply press **<enter>** at this prompt.

- c. The following prompt displays:

Do you wish to configure the interface group for this Virtual LAN at this time? (y)

You can assign physical ports to the new Group at this time. To begin assigning ports to the new Group, press **<Enter>** and go to Step 4.

To assign ports to the Group later, type **n** and **<Enter>**. The new Group is configured but does not yet contain any ports. You can use the **addvp** command later to assign ports to the Group (see *Adding Virtual Ports* on page 17-46). A message similar to the following displays confirming the creation of the new Group.

GROUP 6 has been added to the system.

You may add interfaces to this group using the addvp command at a later date.

For now, the GROUP is inactive until you add interfaces.

Configure Auto-Activated LANE service ? [y/n] (y) :

If you want to configure switch ports later (or simply rely on the dynamic port assignment capability's of the mobile group) skip ahead to *Step 5. Configure Auto-Activated LANE Ports (Mobile Groups Only)* on page 17-33.

Step 4. Configuring Virtual Ports

You can now enter configuration parameters for each switch port to be included in this Group. These configuration parameters include the bridging mode, output format type, and administrative state. In addition, if the port you are configuring is Ethernet (10 Mbps), you can also configure port mirroring.

Prompts for configuring virtual ports follow directly after Group Mobility prompts. You can choose to add ports now or add them later through the **addvp** command. Follow these steps:

- a. After you have stepped through the Routing and/or Group Mobility prompts, the following message displays:

Do you wish to configure the interface group for this Virtual LAN at this time? (y)

You can assign physical ports to the new Group at this time. To begin assigning ports to the new Group, press **<Enter>** and go to Step b.

To assign ports to the Group later, type **n** and **<Enter>**. The new Group is configured but does not yet contain any ports. You can use the **addvp** command later to assign ports to the Group (see *Adding Virtual Ports* on page 17-46). A message similar to the following displays confirming the creation of the new Group.

**GROUP 6 has been added to the system.
You may add interfaces to this group using the addvp command at a later date.
For now, the GROUP is inactive until you add interfaces.**

- b. After indicating that you want to set up ports, the following prompt displays:

**Initial Vports (Slot/Phys Intf. Range) - For example, first I/O Module
(slot 2), second interface would be 2/2. Specify a range of interfaces
and/or a list as in: 2/1-3, 3/3, 3/5, 4/6-8**

Enter the port or ports that you want to include in this new Group. The notation for adding a port to a group is

<slot number of module>/<port number on the module>

The OmniStack's front panel is divided into several areas labeled S1, S2, S3, etc. These areas relate to the conceptual division of the switch into several modules. S1 is the management module (referred to as the MPM), S2 is the uplink module (if the switch supports an uplink module), and S3, S4, etc. are the Ethernet device connection modules.

You may enter multiple ports from multiple switching modules. For example, to add ports 1 through 3 on the module in slot 2, specify **2/1-3**. To additionally add the third and fifth port on the module in the third slot, specify **3/3, 3/5**. The complete slot port specification would be:

2/1-3, 3/3, 3/5

- c. If you enter a port that is already assigned to another Group, then you will be prompted on whether or not you want to change its assignment. A message similar to the following displays for each port that you enter:

**Initial Slot/Interface Assignments: 2/8
2/8 - This interface has already been assigned to GROUP 1 -
(Default GROUP #1).
Do you wish to remove it from that GROUP and assign it (with
new configuration values) to this GROUP (n)?**

Simply enter a **y** at each port prompt to change its Group assignment and begin setting port parameters. You could also enter a **c** at this prompt to accept all default port parameters and skip port configuration prompts. If you enter a **c**, *all* remaining ports are automatically added to the Group with default settings, and your work is complete.

- d. The virtual port configuration menu displays:

```

                                Modify Ether/8 Vport 2/8 Configuration

1) Vport                        : 9
2) Description                  :
3) Bridge Mode                  : Auto-Switched
   31) Switch Timer             : 60
4) Flood Limit                  : 192000
5) Output Format Type           : Default (IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status   : Enabled, inactive
8) Mirrored Port Status        : Disabled, available

Command {Item=Value/?/Help/Quit/Redraw/Next/Previous/Save} (Redraw) :
```

Descriptions for each of the fields in this display follow. To change any default value, enter the line number for item, an equal sign (=), and then the value for the parameter. When you have completed the configuration for this port, enter **next** to begin configuring the next port. Enter **save** to save all configured settings and move onto the next step in the group creation process.

1) Vport

The virtual port number for this port. The next virtual port number available in the switch is shown by default in this field.

2) Description

Enter a useful description for this virtual port using alphanumeric characters. The description may be up to 30 characters long.

3) Bridge Mode

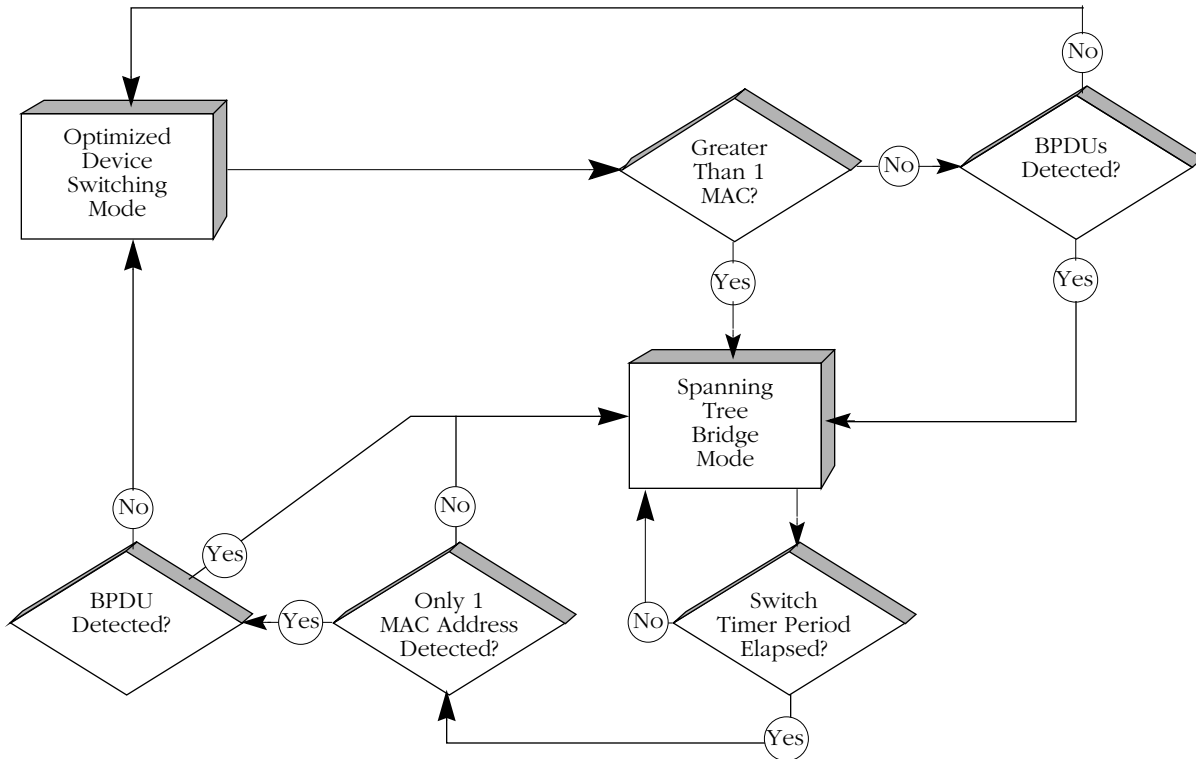
Select the bridge mode used by this port. The choices are:

Spanning Tree Bridge. The default setting for all non-Ethernet ports. This mode is appropriate for backbone and hub connections. The port acts as a standard 802.1d bridge port. It forwards BPDU frames out the port. When frames are received, Spanning Tree BPDUs are processed, and Spanning Tree dynamically controls the forwarding state. If flooding occurs, all frames destined for unknown MAC addresses, broadcast addresses, or multicast addresses will be sent to all ports in the same Group. Enter **3=b** and press **<Enter>** to select Spanning Tree Bridge mode.

Optimized Device Switching. This mode is appropriate for dedicated connections to a single workstation or server. Spanning Tree is turned off. No Spanning Tree BPDUs will be sent and the port will always be in the forwarding state. The port will stay in this mode even if a Spanning Tree BPDU is detected. In addition, all MACs learned will not be aged out (regardless of the Bridge Aging Timer setting) until the port is disconnected or configured to be administratively down. No flooding of packets with an unknown destination address is allowed after at least one MAC address has been learned. (An exception to this rule occurs on Ethernet ports. When these ports are in optimized mode, packets with unknown destination addresses will be flooded.) Packets with a broadcast or multicast destination will always be allowed. Enter **3=o** and press **<Enter>** to select Optimized Device Switching mode.

Auto-Switch. The default setting for all Ethernet ports. This mode is appropriate for dedicated connections requiring a switch-over to bridge mode when multiple devices are detected. A port in Auto-Switch mode will start in Optimized Device Switching mode (see description above). The port will remain in Optimized Device Switching mode until a Spanning Tree BPDU is detected or more than one MAC address transmits data. Once either of these conditions is met, the port will switch to Spanning Tree Bridge mode and Spanning Tree will start (if configured in the switch).

An Auto-Switch port will remain in Spanning Tree Bridge mode as long as there are BPDUs and multiple MACs. However, the port can revert back to Optimized Device Switching Mode if the time specified in the next field (**Switch Timer**) transpires without BPDUs and multiple MACs. Also, if the port is disconnected or configured to be administratively down, then an Auto-Switch port will revert back to Optimized Device Switching mode when it becomes operational again. Enter **3=a** and press **<Enter>** to select Auto-Switch mode.



How Auto-Switch Bridge Mode Works

31) Switch Timer

If you selected the Auto-Switch bridge mode, then you can configure this field. Enter the time-out period, in seconds, for an Auto-Switch port that has turned to Spanning Tree Bridge mode port to revert back to Optimized Switching mode. When in Auto-Switch mode, a port switches to Spanning Tree Bridge mode as soon as it detects a BPDU or more than one MAC address. The port will switch back to Optimized Switching mode after the time-out value you define here.

4) Flood Limit

The flood limit allows you to tune a virtual port to limit the flooding of broadcast, multi-cast, and unknown destination packets. This feature is useful for controlling broadcast storms on your network. While each network is different, in general the amount of flooded traffic represents a relatively small percentage of network traffic.

The flood limit is actually a “transmit credit” that is issued every five (5) seconds. When a packet is flooded on this port, the size of the packet, in bytes, is decremented from the current credit value. The credit value is the value you enter in this field multiplied by five. An additional credit, in the amount of the value you enter here multiplied by five, is allocated to each virtual port every five (5) seconds. If the credit value ever falls below zero, then all flooded packets are discarded until another credit is allocated. Flood limit checking is disabled if you enter a flood limit of zero (0). The flood limit default is 192,000 bytes per second, which equates to a transmit credit of 960,000 bytes every five seconds.

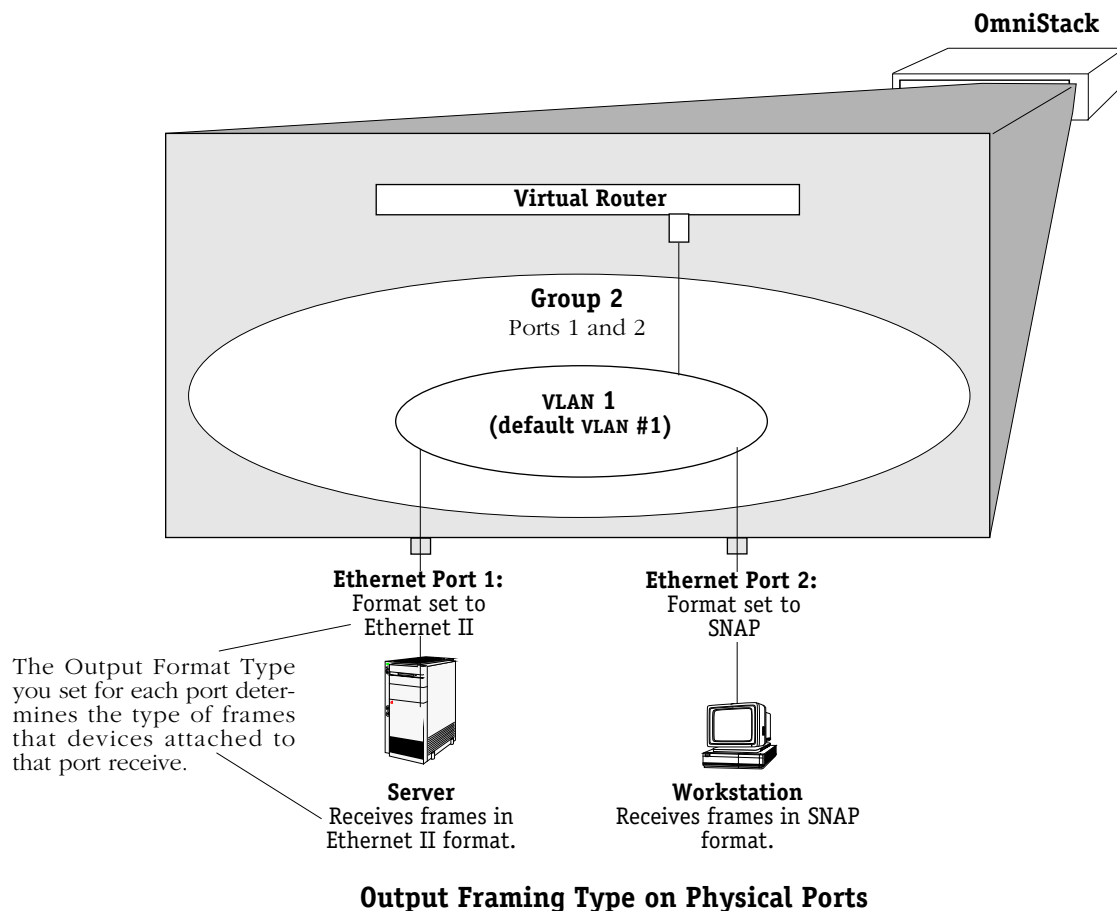
5) Output Format Type

The options will be different for Ethernet, Token Ring, and FDDI ports (note that OmniStack switches do not provide Token Ring or FDDI ports).

The output format setting determines the kind of frame that will be sent out this physical port. If translation is necessary, then incoming frames will be translated to this format before being sent out this port. For example, on an Ethernet port incoming FDDI frames need to be translated to Ethernet. However, there are four types of Ethernet frames—Ethernet II, IPX 802.3, SNAP, and LLC. The format type you select here would determine the frame format to which non-Ethernet frames would be translated. The following figure illustrates how a port's framing type affects communication with attached devices.

Note

This parameter differs from the router framing type selected during the configuration of the virtual router port. The router framing type is the encapsulation done on a router port, whereas this output format type applies only to translations on this virtual port.



Note that for Ethernet, the default output format option is Ethernet II for IP frames and 802.3 for IPX frames. Although the UI provides output options for Token Ring and FDDI, OmniStack switches do not provide Token Ring or FDDI ports.

You can customize your frame translation settings even further through the Switch menu. The Switch menu allows you to set translations at the frame format level (i.e., incoming SNAP frames could be translated one way, while incoming LLC frames could be translated another way) based on protocol type (IP or IPX). The Switch menu is explained in Chapter 16, “Configuring LAN Switch Translations.”

6) Ethernet 802.2 Pass Through

For Ethernet ports only. If you answer **Yes** to this prompt, then frames received in the IEEE 802.2 format will not be translated according the Output Format Type chosen in line 5; they will be sent as is in their native IEEE 802.2 format. If you answer **No**, then 802.2 frames will be subject to the Output Format Type chosen in line 5.

7) Admin, Operational Status

Select whether to administratively enable or disable this port. When you enable the port, the port can transmit and receive data as long as a cable is connected and no physical or operational problems exist. When you disable a port, the port will not transmit or receive data even if a cable is connected and the physical connection is operational. If you disable the port at this point, you can enable it later through the **modvp** command (see *Modifying a Virtual Port* on page 17-47).

8) Mirrored Port Status

If the port you are configuring is Ethernet (10 or 10/100 Mbps), you can set up port mirroring. You can mirror traffic on this port to another like port. Port mirroring is a useful feature for monitoring traffic on particular ports. It is discussed in more detail later in this chapter in *Port Mirroring* on page 17-58.

If you want to mirror this port, enter a **8=e**, press **<Enter>** and you will be prompted for the slot and port number of the “mirroring” port (i.e., the port that can “see” all traffic for this port):

Mirroring vport slot/port ? () :

Enter the mirroring port’s slot and port number and press **<Enter>**.

If port mirroring is not supported on this port, then the following prompt will display:

mirroring not supported on this port type

After the port mirroring prompts, the switch confirms the addition of the port to the group with a message similar to the following:

Adding port 2/8 to Group 6. . .

Make configuration changes to the port until you are satisfied. Use the next command to configure all remaining virtual ports that you need to configure. If you have completed the final virtual port, then your work is complete. You can always alter Group parameters (including virtual router parameters for the default VLAN) later through the **modvl** command (see *Modifying a Group or VLAN* on page 17-42) and modify virtual port parameters through the **modvp** command (see *Modifying a Virtual Port* on page 17-47).

Step 5. Configure Auto-Activated LANE Ports (Mobile Groups Only)

If you set up a mobile group in Step 3, you can configure LANE service ports as members of this group. You can automatically insert these services into a mobile group without any more configuration than the following steps.

- a. After you have completed configuring virtual ports, the following prompt displays:

Configure Auto-Activated LEC service ? [y/n] (y) :

To add a LANE service to this mobile group, enter a **Y** and press **<enter>** at this prompt.

- b. The following prompt displays:

Enter the primary slot/interface for this service:

Enter the slot and port of the ATM access switch port that will be the primary port for this LANE service. The port should be an ATM access port as opposed to an Ethernet port.

- c. The following prompt displays

Enter the secondary slot/interface for this service (hit enter for none):

Enter the slot and port of the ATM access switch port that will be used as the secondary port for this LANE service. This port will be used in case the primary port is not available; it will remain dormant until the primary port goes down. If you do not want to specify a secondary port (or if there is no other ATM access port in the switch), press **<Enter>**.

- d. The following prompt displays

Enter E-LAN name:

Enter the name you want to use for this LANE service. Duplicate ELAN names cannot be used on the same port. It is possible to set up LANE services through the **cas** command, so make sure you are not using an ELAN name that has already been used. If you create a service through **cas** that uses the same name as a dormant auto-activated service, then the auto-activated service will not become operational.

- e. The following prompt displays

Enter the ELAN Type [Ethernet|Token Ring] (Ethernet):

Ethernet and Token Ring LANE clients are handled differently by switch software. Specify which type you want to configure. Note that OmniStack switches do not provide Token Ring ports, however they do support Token Ring LANE clients.

- f. The following prompt displays

Enable this service ? [y/n] (y):

Indicate whether or not you want to enable this service. If you enable this service it will remain operationally inactive until a device becomes a member of this ELAN. As soon as you enable the service, devices will be able register with this ELAN and become members of this mobile group. Disabling the service is a good choice if you are still configuring your network and want to test it further.

- g. The following prompt displays

Configure another service for this group ? [y/n] (n):

If you want to configure more auto-activated LANE service ports, enter a **Y** at this prompt. To begin configuring AutoTracker rules for this mobile group, enter an **N** and continue with Step 6.

Step 6. Configuring AutoTracker Policies (Mobile Groups Only)

When you have completed configuring mobile group and auto-activated LANE services, you can begin configuring AutoTracker policies for this mobile group. Instructions for configuring these rules can be found in Chapter 18, “Configuring Group and VLAN Policies.” Please refer to that chapter for instructions on configuring each policy type. After you configure AutoTracker policies, you are done configuring this mobile group and a prompt similar to the following displays:

VLAN 9: 1 created successfully

You can configure rules for this group later through the **modatvl** command. This command also works with mobile groups as long as you indicate you want to alter VLAN 1 in the mobile group (i.e., the command line would read **modatvl 3:1** to modify mobile group 3).

Creating a WAN Routing Group

After entering basic Group information as described in *Step 1. Entering Basic Group Information* on page 17-21, you should have answered Yes to the following prompt:

Enable WAN Routing? (n):

if you want to enable WAN Routing. WAN Routing Groups are treated differently than other Groups, as described earlier. The following steps complete the configuration of the WAN Routing Group.

- a. After answering **y** to the **Enable WAN Routing?** prompt, the following prompt displays:

Enable IP (y):

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

Note

You may enable routing of both IP and IPX traffic over a WAN connection. If you set up dual-protocol routing, you must fill out information for both IP and IPX parameters.

- b. 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 of the default VLAN within this Group. After you enter the address, press **<Enter>**.

- c. The following prompt displays:

IP Subnet Mask (0xfffff00):

The default IP subnet mask (in parentheses) is automatically derived from the default VLAN 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>**.

- d. The following prompt displays:

IP Broadcast Address (198.200.10.255):

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

- e. 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>**.

- f. 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 the default VLAN in this Group 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 the default VLAN in this Group 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. The default VLAN in this Group 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. The default VLAN in this Group 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. The default VLAN in this Group will neither send nor receive routing information to/from other VLANs. Enter **i** and press **<Enter>** to select Inactive mode.

- g. 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 the default VLAN in this WAN Group will not be able to 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 for the default VLAN in this Group, enter **n** and press **<Enter>**. You can always set up IPX routing for other VLANs within this Group.

You are done configuring this WAN Routing Group. See the appropriate WAN interface chapter for further information on configuring this Routing service.

- h. 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 digits in this field. If you enter less than eight hex digits, the system prefixes your entry with zeros to create eight digits.

- i. 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>**.

- j. 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)}
RIP and SAP triggered (t)}          (a):
  
```

Select how you want the IPX protocols, RIP (router internet protocol) and SAP (service access protocol), to be configured for the default VLAN in this Group. 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 default 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 default 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 default VLAN to which it is attached does not participate in either RIP nor SAP updates. Enter an **i** and press **<Enter>** to select RIP and SAP inactive.

RIP and SAP triggered. The IPX router port is active, but RIP and SAP information will be sent out on the port only when a network change has occurred. This option is more cost effective for WAN links and is best suited for smaller network environments that don't change often. Enter a **t** and press **<Enter>** to select RIP and SAP triggered.

When you are done entering Router parameters, a message similar to the following displays:

```

GROUP 5 has been added to the system
  
```

Creating an ATM CIP Group

After entering basic Group information as described in *Step 1. Entering Basic Group Information* on page 17-21, you should have answered **Yes** to the following prompt:

Enable ATM CIP? (n):

if you want this Group to support ATM Classical IP (CIP). ATM CIP Groups are different from other Groups. CIP Groups must be created if there are devices in the ATM network that support only Classical IP or require CIP to communicate over ATM. The following steps complete the configuration of the ATM CIP Group.

- a. After answering **y** to the **Enable ATM CIP?** prompt, 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 of the default VLAN within this Group. After you enter the address, press **<Enter>**.

- b. The following prompt displays:

IP Subnet Mask (0xffffffff):

The default IP subnet mask (in parentheses) is automatically derived from the default VLAN 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>**.

- c. The following prompt displays:

IP Broadcast Address (198.200.10.255):

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

- d. The following prompt displays:

Description (30 chars max):

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

- e. 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 the default VLAN in this Group 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 the default VLAN in this Group 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. The default VLAN in this Group 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. The default VLAN in this Group 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. The default VLAN in this Group will neither send nor receive routing information to/from other VLANs. Enter **i** and press **<Enter>** to select Inactive mode.

When you are done entering Router parameters, a message similar to the following displays:

GROUP 6 has been added to the system

You should now follow the instructions for configuring an ATM Classical IP Service described in Chapter 28, "Configuring ATM Services."

IPX Routing Over ATM

IPX routing is not supported directly over ATM. IPX can be routed by the OmniStack and switched out one of the ATM services. See Chapter 28, "Configuring ATM Services," for more information.

Viewing Current Groups

The **gp** command provides information on all currently defined Groups in a switch including Group number, network address, protocol type, and encapsulation type. You can obtain information on all groups in a switch by entering:

```
gp
```

A screen similar to the following displays:

Group ID (:VLAN ID)	Group Description	Network Address (IP Subnet Mask) or (IPX Node Addr)	Proto/ Encaps
=====	=====	=====	=====
1	Default GROUP (#1)	198.206.182.115 (ff.ff.ff.00)	IP / ETH2
2	New GROUP (#2)	198.206.101.12 (ff.ff.ff.00)	IP / SNAP
3	New GROUP (#3)	198.206.181.10 (ff.ff.ff.00)	IP / 1490
4	New Group (#4)	198.206.183.44 (ff.ff.ff.00) 12314526 (0020da:020484)	IP / ETH2 IPX / 8023
5	New GROUP	198.206.143.11 (ff.ff.ff.00)	CIP / 1483

You can also get information on a specific Group by entering **gp** followed by the Group number. For example,

```
gp 3
```

displays information just on Group 3:

Group ID (:VLAN ID)	Group Description	Network Address (IP Subnet Mask) or (IPX Node Addr)	Proto/ Encaps
=====	=====	=====	=====
3	New GROUP (#3)	198.206.181.10 (ff.ff.ff.00)	IP / 1490

The following sections describe the columns in this table:

Group ID (:VLAN ID). The identification number assigned to this Group when it was created through the **crgp** command. The Group identifier is typically consistent network-wide (i.e., Group 3 in this switch should be the same Group as Group 3 configured in all other OmniStackes in the network). If this Group contains any VLANs, then they will be listed below the Group number. If the default VLAN in the Group supports both IP and IPX routing, then information on both (network address, etc) will display. Group 4 in the screen sample above shows a case where both IP and IPX routing are supported.

Group Description. The textual description of this Group that was entered when the Group was created or modified. This description is limited to 30 characters.

Network Address (IP Subnet Mask) or (IPX Node Addr). For each virtual router port configured, two addresses are listed. Both of these addresses were configured when the Group was created or modified through **crgp** or **modvl**. The first address is the Network Address, which is the address of the virtual router port for the default VLAN (VLAN #1) in this Group. For an IP virtual router port, this address is the IP address, which is shown in dotted decimal format. For an IPX virtual router port, this address is the IPX network address, which is shown as eight hex characters.

A second address is displayed below the Network address. For IP, this address is the IP Subnet Mask, which is normally derived from the default VLAN IP address class. For IPX, this address is the IPX Node Address.

Proto/Encaps. For each Group or VLAN listed, the top field is the Protocol supported by this virtual router port. Possible values in the field are: **IP** (IP router), **IPX** (IPX router), and **CIP** (Classical IP Group with CIP router). If you configured an IP and an IPX router port, then two router entries will be listed—one with a Protocol of IP and the other with a Protocol of IPX.

The bottom field is the encapsulation used for outgoing frames on the router port. This encapsulation was configured when the router port was configured. Possible values for this field depend on the Protocol and type of Group.

Frame Relay WAN Groups will always display **1490** to indicate RFC 1490 encapsulation is performed on frames. ATM Classical IP (CIP) Groups will display **1483** to indicate RFC 1483 encapsulation is performed on frames.

IP and IPX routers have additional possible encapsulation types. For IP virtual router ports, the possible encapsulation types are as follows:

- **ETH2** Ethernet II
- **SNAP** Ethernet 802.3 SNAP
- **FDDI** FDDI
- **8025** Token Ring 802.5
- **TSRS** Token Ring Source Routing

For IPX virtual router ports, the possible encapsulation types are as follows:

- **ETH2** Ethernet II
- **LLC** Ethernet 802.3 LLC
- **SNAP** Ethernet 802.3 SNAP
- **8023** Ethernet 802.3 (Novell raw)
- **FDDI** FDDI SNAP
- **FSRS** FDDI Source Routing SNAP
- **FLLC** FDDI LLC
- **FSRL** FDDI Source Routing LLC
- **8025** Token Ring SNAP
- **TSRS** Token Ring Source Routing SNAP
- **TLLC** Token Ring LLC
- **TSRL** Token Ring Source Routing LLC

Modifying a Group or VLAN

After creating a Group (through **crgp**) or VLAN (through **cratvl**, see Chapters 17 and 18), you can change any of their parameters through the **modvl** command. In addition, if you did not set up a virtual router port (IP or IPX) during the initial Group or VLAN configuration, you can set one up with **modvl**. To use this command, enter **modvl** followed by the Group number and VLAN number to change. For example, to modify parameters in Group 2, VLAN 1, enter:

```
modvl 2
```

Note that you do not need to specify a VLAN number to modify the default VLAN within a Group. To modify parameters in Group 2, VLAN 2, you would enter:

```
modvl 2:2
```

A screen similar to the following displays.

Current values associated with GROUP 2.1 are as follows:

```
1) GROUP Number      - 2:1
2) Description        - New GROUP (#2)
IP Parameters:
3) IP enabled         - Y
4) IP Network Address - 198.206.101.12
5) IP Subnet Mask     - 255.255.255.0
6) IP Broadcast Address - 198.206.101.255
7) Router Description - Router Port #2
8) RIP Mode           - Silent
   {Active (a), Inactive (i), Deaf (d), Silent (s)}
9) Routing disabled   - N
10) NHRP enabled      - N
11) Default Framing   - Ethernet II
   {Ethernet II(e), Ethernet 802.3 (8), fddi (f),
   token ring (t), source route token ring (s)}
IPX parameters:
12) IPX enabled       - N

(save/quit/cancel)
:
```

The Group number at the top of this sample screen is followed by the number 1 (**GROUP 2.1**), meaning that the information applies to default VLAN #1 in this Group. If this screen displayed information on Group 2, VLAN 2, then this field would read **GROUP 2:2**.

The colon prompt (:) at the bottom of the screen is used to prompt for user input. To change a value, type the line number of the item you want to change, followed by an equal sign (=) and the new value. For example, to set a new description you could enter:

```
2=Engineering
```

All of the **modvl** parameters are described in the section for creating a new Group, *Creating a New Group* on page 17-20.

Note

Line numbering for the **modvl** command will vary depending on whether you have an IP or IPX router configured. Each type of router contains several parameters that require extra line numbers.

Viewing Your Changes

When you enter a change at the colon prompt, the **modvl** screen does not normally refresh. If you want to see the current Group or VLAN settings, including any changes you made, enter a question mark (?) at the colon prompt. The **modvl** screen will refresh.

Saving Your Changes

Once you have entered all your modifications and you want to save them, type **save** at the colon prompt. You will exit the **modvl** command and your changes will take effect.

Canceling Your Changes

You can also exit the **modvl** command without saving any changes you made in the current session. Simply enter **cancel** at the colon prompt or enter **<Ctrl>-d**. The **modvl** command will end and none of the changes you made will be saved.

Changing the IP Address

Changing the IP address can also affect the Subnet Mask and the Broadcast Address. The new IP address means that the Subnet Mask and Broadcast Address must be re-generated and the following message displays:

New IP address generates new subnet and broadcast address
Enter '?' to view the changes

The system automatically creates new Subnet Mask and Broadcast addresses based on the new IP address. If you enter a question mark (?) at this point you could view these changes.

If you remove the last IP address in the system, you will see a warning message that SNMP (and other applications) are now inoperational.

Changing the IP Subnet Mask

Changing the IP Subnet Mask can also affect the IP Broadcast Address. The new Subnet Mask means that the Broadcast Address must be re-generated and the following message displays:

New mask caused change in broadcast address

The system automatically created a new Broadcast address based on the new Subnet Mask. If you entered a question mark (?) at this point you could view these changes.

Enabling IP or IPX Routing

If you enable IP or IPX routing by setting the corresponding **modvl** lines from **N** to **Y**, then the screen automatically refreshes with additional lines for the new router port parameters. All lines are set to router defaults. The router defaults are as follows:

IP Router

IP Network Address	0.0.0.0
IP Subnet Mask	0.0.0.0
IP Broadcast Address	0.0.0.0
Router Description	(no description shown for default)
RIP Mode	Silent
Default Framing Type	Ethernet II

IPX Router

IPX Network Address	0x0
Router Description	(no description shown for default)
RIP/SAP Mode	RIP and SAP are active
Default Framing Type	Ethernet II

You can change any of these defaults as you would any other **modvl** parameters: enter the line number, followed by an equal sign (=) and the new parameter.

Note

You must at least enter a Network Address for a new router or you will not be able to save the configuration.

Deleting a Group

You can delete a Group as long as it does not contain any virtual ports. The default Group, Group #1, cannot be deleted. To delete a Group, enter **rmgp** followed by the Group number you want to delete. For example, if you wanted to delete Group 5, you would enter:

```
rmgp 5
```

If the Group does not contain any virtual ports, then a confirmation message displays:

```
GROUP 5 removed.
```

If the Group still contains virtual ports, then a message similar to the following displays:

```
GROUP 5 has active entries, you must remove  
these prior to removing the GROUP (use rmvp for this).
```

You must first remove the Group's virtual ports before the Group can be removed. The **rmvp** command allows you to remove virtual ports. See *Deleting a Virtual Port* on page 17-48 for information on using this command.

Note

Some commands in the Bridge Management menu (described in Chapter 15, "Configuring Bridging Parameters") require you to select a Group before making configuration changes. If you delete the currently selected Group with **rmgp**, then the new currently selected Group reverts to the default Group, Group #1.

Adding Virtual Ports

You can add virtual ports to a Group at any time after the Group is created. The **addvp** command allows you to add one or more ports to a Group you specify. If you have used the **crgp** command to add virtual ports, then you will find the **addvp** command fields very familiar.

To use **addvp**, enter the command followed by the Group number to which you want to add the port. Next, specify the port or ports you want to add.

addvp <Group Number for port> <Module Slot>/<Port Number>

For example, if you wanted to add ports 4 through 6 on the module in slot 4 to Group #5, then you would specify:

addvp 5 4/4-6

The procedure for using **addvp** is as follows:

1. Enter **addvp** followed by the Group number where you want this port to reside, followed by the physical slot and port numbers you want to configure.
2. If you enter a port that is already assigned to another Group, then you will be prompted on whether or not you want to change its assignment. A message similar to the following displays for each port that you enter:

**4/4 - This interface has already been assigned to GROUP 1 -
(Default GROUP #1).
Do you wish to remove it from that GROUP and assign it (with
new configuration values) to this GROUP (n)?**

Simply enter a **y** at each port prompt to change its Group assignment and begin setting port parameters. You could also enter a **c** at this prompt to accept all default port parameters and skip port configuration questions. If you enter a **c**, *all* remaining ports are automatically added to the Group with default settings, and your work is complete.

3. The virtual port configuration menu displays:

Modify Ether/8 Vport 4/4 Configuration

```

1) Vport                : 9
2) Description          :
3) Bridge Mode          : Auto-Switched
   31) Switch Timer     : 60
4) Flood Limit          : 192000
5) Output Format Type   : Default (IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status : Enabled, inactive
8) Mirrored Port Status  : Disabled, available
  
```

Command {Item=Value/?/Help/Quit/Redraw/Next/Previous/Save} (Redraw) :

Descriptions for each of the fields in this display begin on page 17-29. To change any default value, enter the line number for the item, an equal sign (=), and then the value for the parameter. When you have completed the configuration for this port, enter **save** to save all configured settings.

Modifying a Virtual Port

You can modify a virtual port through the **modvp** command. The **modvp** command is very similar to the **addvp** command and the port configuration phase of the **crgp** command. To use **modvp**, enter the command, followed by the Group number for the port, and the physical slot and port number for the port:

modvp <Group Number for port> <Module Slot>/<Port Number>

You can specify only one port at a time. For example, if you wanted to modify the parameters for Port 7 on the module in Slot 4, and the Port currently resides in Group 6, then you would enter:

modvp 6 4/7

The procedure for using **modvp** is as follows:

1. Enter **modvp** followed by the Group number where the port currently resides, the physical slot and port number.
2. A prompt displays requesting your confirmation:

Modify local port 7 (Virtual port (#14)) ? (y) :

Simply press **<Enter>** if this is the correct virtual port. The Virtual Port number in parentheses (**Virtual Port #14** in this case) is the virtual port number within this entire OmniStack. Virtual ports are numbered sequentially within the switch, not within a Group or VLAN.

3. The virtual port configuration menu displays:

Modify Ether/8 Vport 4/7 Configuration

```

1) Vport                : 9
2) Description          :
3) Bridge Mode          : Auto-Switched
   31) Switch Timer     : 60
4) Flood Limit          : 192000
5) Output Format Type   : Default (IP-Eth II, IPX-802.3)
6) Ethernet 802.2 Pass Through : Yes
7) Admin, Operational Status : Enabled, inactive
8) Mirrored Port Status  : Disabled, available

```

Command {Item=Value/?/Help/Quit/Redraw/Next/Previous/Save} (Redraw) :

Descriptions for each of the fields in this display begin on page 17-29. To change any default value, enter the line number for the item, an equal sign (=), and then the value for the parameter. When you have completed the configuration for this port, enter **save** to save all configured settings.

Deleting a Virtual Port

You can delete a virtual port from its existing Group by using the **rmvp** command. When you remove a virtual port, the port is moved to the default switch Group, Group #1, and all port parameters are reset to defaults except for the port name. For example, if you configured a port with a special flood limit and customized translation settings and you then removed the port, you would lose those port settings.

To remove a port, enter the **rmvp** command, followed by the Group number where the port currently resides and the physical slot and port number for the port:

rmvp <Group number> <Module Slot>/<Port Number>

For example, to delete Port 7 on the module in Slot 4, and the Port currently resides in Group 6, you would enter:

rmvp 6 4/7

A prompt displays requesting that you confirm the deletion:

Local port 7 (Virtual po...) is attached to this slot/interface - remove? (n):

Enter a **y** and press **<Enter>** to remove the port. Another message displays confirming the deletion:

BRIDGE port on 4/7 moved to GROUP 1.

If the port you specified did not exist in the Group you specified in the **rmvp** command, then a message similar to the following would display:

Specified port(s) not found on GROUP 6.

Viewing Information on Ports in a Group

The **via** command allows you to view port attachments associated with a specified Group or all Groups in a switch. Entering

```
via
```

displays summary information for all virtual ports in the switch. You can also display virtual interface attachments for a specific Group by specifying the Group ID after the **via** command. For example, to view ports for Group 2, you would enter

```
via 2
```

The same type of information is displayed for a single Group as is displayed for all Groups. The following screen shows a sample from the **via** command when specified without a Group ID.

GROUP Interface Attachments For All Interfaces

GROUP: Slot/Intf	Description	Service/ Instance	Protocol	Admin Status
=====	=====	=====	=====	=====
1.1 : *	GROUP #1.0 IP router vport	Rtr / 1	IP	Enabled
2.1 : *	for group 2	Rtr / 2	IP	Enabled
1:2/1	Virtual port (#2)	Brg / 1	Tns	Enabled
1:2/2	Virtual port (#3)	Brg / 1	Tns	Enabled
1:2/3	Virtual port (#4)	Brg / 1	Tns	Enabled
2:2/4	finance server	Brg / 1	Tns	Enabled
1:2/5	Virtual port (#6)	Brg / 1	Tns	Enabled
1:2/6	Virtual port (#7)	Brg / 1	Tns	Enabled
1:2/7	Virtual port (#8)	Brg / 1	Tns	Enabled
1:2/8	Virtual port (#9)	Brg / 1	Tns	Enabled
1:3/1	Virtual port (#1)	Brg / 1	Tns	Enabled
1:4/1	Virtual port (#10)	Brg / 1	Tns	Enabled
1:4/2	Virtual port (#11)	Brg / 1	Tns	Enabled
1:4/3	Virtual port (#12)	Brg / 1	Tns	Enabled
1:4/4	Virtual port (#13)	Brg / 1	Tns	Enabled
1:4/5	Virtual port (#14)	Brg / 1	Tns	Enabled
1:4/6	Virtual port (#15)	Brg / 1	Tns	Enabled

GROUP: Slot/Intf. **GROUP** is the group number to which this port is assigned. When the Group displays as a Group number followed by a decimal and a 1 (**1.1** and **2.1** in the above sample), it represents the router port on the default VLAN within that Group. **Slot** is the position in the chassis of the switching module where this port is located. **Intf** (Interface) is the physical port on the switching module. When the Slot and Interface are shown as an asterisk (*)—as the top two entries in the above table display—it represents a virtual router port that does not have a corresponding physical interface.

Description. The textual description entered for either the virtual router port or the virtual switch port. This description was entered through **crpg** or **modvl** for virtual router ports, or through **crgp**, **addvp**, or **modvp** for virtual switch ports.

Service/Instance. **Service** is the service type configured for this port. **Instance** is an identifier of this service type within the switch. For example, multiple virtual router ports within the switch will be labelled consecutively (1, 2, 3, etc.), and will each have a different **Instance** number.

Values for the service type are as follows (note that OmniStack switches do not provide FDDI or Frame Relay ports):

- **Rtr** Virtual router port
- **Brg** Virtual bridge port
- **Tnk** Virtual trunk port (used for ATM, FDDI, and WAN)
- **T10** 802.10 FDDI service port
- **FRT** Frame Relay trunk port
- **Lne** LAN Emulation service port
- **CIP** Classical IP service port
- **Vlc** VLAN Clusters (X-LANE) service port

Protocol. The bridging protocol for virtual ports and services or the routing protocol for virtual router ports. Possible values are (note that OmniStack switches do not provide FDDI or Frame Relay ports):

- **Tns** Transparent bridge. Bridges maintain a dynamic table of known MAC addresses on connected segments. The table is used to make forwarding decisions. When a frame is received that contains a destination address that matches an address in the table, it is forwarded to designated bridge ports that are in forwarding state.
- **SR** Source Routing Bridge. Normally used in Token Ring environments. Routing information is determined by looking at the Routing Information Field (RIF) in a frame. The RIF contains the segment and bridge numbers that create the path to the destination. Virtual ports are configured as Source Routing bridges through the **src** command.
- **SRT** Source Routing Transparent. Normally used in Token Ring environments. Allows Source Routing and Transparent bridges to coexist. The Source Routing Transparent Bridge will form a Spanning Tree with other Transparent Bridges and Source Routing Transparent Bridges and will forward frames that do not contain a Routing Information Field (RIF) to destinations reachable by the Spanning Tree. If the bridge detects routing information in the RIF, it will forward it the same way Source Routing bridges do. Virtual ports are configured as Source Routing Transparent bridges through the **src** command.
- **IP** IP Routing Protocol. Routing Information Protocol (RIP) used to learn routes from neighboring routers. You configure an IP router through the **crpg** or **modvl** commands. Other IP routing parameters can be set through the Networking menu commands, which are described in Chapter 23, "IP Routing."
- **IPX** IPX Routing Protocol. Uses RIP to learn routes from neighboring routers and the Service Advertising Protocol (SAP) to maintain a database of network services for requesting workstations. Other IPX routing parameters can be set through the Networking menu commands, which are described in Chapter 25, "IPX Routing."
- **CIP** Classical IP Routing (RFC 1577). Classical IP is necessary when an ATM network contains devices that support only CIP. This type of routing is configured when you initially create a Group through the **crpg** command.
- **FR** Frame Relay IP Routing. WAN Routing Groups are configured slightly different from other Groups. Frame Relay IP Routing is IP Routing with some enhancements to account for the Frame Relay network.

Admin Status. Indicates whether the port is administratively **Enabled** or **Disabled**. When **Enabled**, the port can transmit and receive data as long as a cable is connected and no physical or operational problems exist. When **Disabled**, the port will not transmit or receive data even if a cable is connected and the physical connection is operational. You can set the Admin Status during port configuration phase of the **crpg**, **addvp**, or **modvp** commands.

Viewing Detailed Information on Ports

The **vi** command displays detailed information about virtual ports. Entering

```
vi
```

displays information for all virtual ports in the switch. You can also display information for only ports in a specific Group by specifying the Group ID after the **vi** command. For example, to view information only for ports in Group 6, you would enter

```
vi 6
```

The same type of information is displayed for a single Group as is displayed for all Groups. The following screen shows a sample from the **vi** command when specified without a Group ID.

Virtual Interface Summary Information- For All Interfaces									
		Slot/ Type/				Status			
Group	Intf	Inst/Srv	MAC Address	Prt	Encp	Admin	Oper	Spn Tr	Mode
1 All	Rtr/	1	0020da:020d40	IP	ETH2	Enabl	Active	N/A	N/A
2 All	Rtr/	2	0020da:020d43	IP	ETH2	Enabl	Active	N/A	N/A
2 All	Rtr/	3	0020da:020d44	IP	ETH2	Enabl	Active	N/A	N/A
1 3/1	Brg/	1/ 1	0020da:048730	Tns	DFLT	Enabl	Inactv	Disabl	Bridged
1 4/1	Brg/	1/ na	0020da:030990	Tns	DFLT	Enabl	Active	Fwdng	Bridged
1 4/2	Brg/	1/ na	0020da:030991	Tns	DFLT	Enabl	Inactv	Disabl	Bridged
1 4/3	Brg/	1/ na	0020da:030992	Tns	DFLT	Enabl	Inactv	Disabl	Bridged
1 4/4	Brg/	1/ na	0020da:030993	Tns	DFLT	Enabl	Inactv	Disabl	Bridged
1 4/5	Brg/	1/ na	0020da:030994	Tns	DFLT	Enabl	Inactv	Disabl	Bridged
1 4/6	Brg/	1/ na	0020da:030995	Tns	DFLT	Enabl	Inactv	Disabl	Bridged
1 4/7	Brg/	1/ na	0020da:030996	Tns	DFLT	Enabl	Inactv	Disabl	Bridged
2 4/8	Brg/	1/ na	0020da:030997	Tns	DFLT	Enabl	Inactv	Disabl	Bridged
1 5/1	Brg/	1/ na	0020da:022860	Tns	DFLT	Enabl	Inactv	Disabl	Bridged

Group. The Group number to which this port is currently assigned.

Slot/Intf. The slot (**Slot**) is the position in the chassis of the switching module where this port is located. The interface (**Intf**) is the physical port on the switching module. If this column reads **All**, then this port is a router port that supports all virtual ports in the Group.

Type/Inst/Srv. The Service Type (**Type**), Instance (**Inst**) of this Service Type in the switch, and service number (**Srv**) for this virtual port. Service Type values are as follows (note that OmniStack switches do not provide FDDI or Frame Relay ports):

- **Rtr** Virtual router port
- **Brg** Virtual bridge port
- **Tnk** Virtual trunk port (used for ATM, FDDI, and WAN)
- **T10** 802.10 FDDI service port
- **FRT** Frame Relay trunk port
- **Lne** LAN Emulation service port
- **Vlc** VLAN clusters (X-LANE) service port
- **CIP** Classical IP service port

The Instance (**Inst**) is an identifier of this type of service within the switch. For example, if more than one virtual router port is configured in the switch, then each “instance” of a router will be given a different number. The service number (**Srv**) is port-specific. If a port has more than one service configured on it, then each service will be identified by a different service number.

MAC Address. The MAC address for this virtual port. Each virtual port is allocated a MAC address.

Prt. The bridging or routing protocol supported by this virtual port. Descriptions of these protocol types are provided on page 17-50. Possible values are (note that OmniStack switches do not provide Frame Relay ports):

- **Tns** Transparent Bridge
- **SR** Source Routing Bridge
- **SRT** Source Routing Transparent Bridge
- **IP** IP Routing Protocol
- **IPX** IPX Routing Protocol
- **CIP** Classical IP Routing (RFC 1577)
- **FR** Frame Relay IP Routing

Encp. Encapsulation used for outgoing packets on this virtual router or switch port. Possible encapsulation values are (note that OmniStack switches do not provide FDDI, Token Ring, or Frame Relay ports):

- **DFLT** Default format for this switch port (differs for each interface type)
- **SWCH** Frame translations have been customized through the Switch menu
- **ETH2** Ethernet II
- **ESNP** Ethernet 802.3 SNAP (virtual router ports)
- **ELLC** Ethernet 802.3 LLC (IPX router ports only)
- **8023** Ethernet 802.3, Novell Raw (IPX router ports only)
- **8025** Token Ring 802.5 SNAP (virtual router ports)
- **TSRS** Token Ring Source Routing SNAP (virtual router ports)
- **TLLC** Token Ring LLC (IPX router ports only)
- **TSRL** Token Ring Source Routing LLC (IPX router ports only)
- **FDDI** FDDI SNAP (virtual router ports)
- **FSRS** FDDI Source Routing SNAP (IPX router ports only)
- **FLLC** FDDI LLC (IPX router ports only)
- **FSRL** FDDI Source Routing LLC (IPX router ports only)
- **1490** Frame Relay Routing (RFC 1490)
- **1483** Classical IP Routing (RFC 1483)
- **SNAP** SNAP (switch ports only)
- **LLC** LLC (switch ports only)

Admin. Indicates whether the port is administratively Enabled or Disabled. When **Enabld**, the port can transmit and receive data as long as a cable is connected and no physical or operational problems exist. When **Disabld**, the port will not transmit or receive data even if a cable is connected and the physical connection is operational. You can set the Administrative Status during the port configuration phase of the **crgp** command, the **addvp** command, or the **modvp** command. A port can have an Administrative Status of Enabled, but still be operationally Inactive. See the description of the **Oper** column below.

Oper. Indicates the current Operational Status of the port. The port will be Active (**Active**) or Inactive (**Inactv**). If the port is Active, then the port can pass data and has a good physical connection. If it is Inactive, then it may not have a good physical connection and it is not capable of passing data at this time.

Spn Tr. The port's current state as defined by the Spanning Tree Protocol. The possible Spanning Tree States are: Disabled, Blocking, Listening, Learning, and Forwarding. This state controls the action a port takes when it receives and transmits a frame. For ports which are Administratively disabled or Operationally Inactive, this state will be Disabled (**Disabl**), meaning the Spanning Tree algorithm is not active on this port. If the state is **Blocking**, then only BPDUs will be transmitted and received. If the state is **Forwarding**, then both data and BPDU frames will be transmitted and received. This Spanning Tree Protocol state is not applicable to virtual router ports and will read **N/A** for those ports.

Mode. The Bridge Mode currently in use on this port. This mode is chosen during the port configuration phase of the **crgp** command, through the **addvp** command, or through the **modvp** command. It is not applicable to virtual router ports and will read **N/A** for those ports. Possible values are:

- **Bridged** Spanning Tree Bridge.
- **AutoSw** Auto Switch.
- **Optimzd** Optimized Device Switching.

See page 17-29 for a description of these bridge modes.

Viewing Port Statistics

The **vs** command displays transmit and receive statistics for ports in the switch. Entering

vs

displays statistics for all virtual ports in the switch. You can also display statistics for only ports in a specific Group by specifying the Group ID after the **vs** command. For example, to view statistics only for ports in Group 6, you would enter

vs 6

You can also display statistics for a specific port by entering the slot and port number after the **vs** command. For example, to view statistics only for Port 1 on the module in Slot 4, you would enter

vs 4/1

The same type of information is displayed for a single Group or port as is displayed for all ports in a switch. The following screen shows a sample from the **vs** command when specified without any Group or port parameters.

Virtual Interface Statistical Information- For All Interfaces

Slot/ Group	Intf	Service/ Instance	Frames In Out	Octets In Out	UcastPkts In Out	NUcastPkts In Out
1 All	Rtr/	1				
2 All	Rtr/	2				
3 All	Rtr/	3				
1 3/1	Tnk/	1	0	0	0	0
			0	0	0	0
1 4/1	Brg/	1	17774	1739560	1707	16067
			684	103048	681	3
1 4/2	Brg/	1	0	0	0	0
			0	0	0	0
1 4/3	Brg/	1	0	0	0	0
			0	0	0	0
1 4/4	Brg/	1	0	0	0	0
			0	0	0	0
1 4/5	Brg/	1	0	0	0	0
			0	0	0	0
1 4/6	Brg/	1	0	0	0	0
			0	0	0	0
1 4/7	Brg/	1	0	0	0	0
			0	0	0	0
1 4/8	Brg/	1	0	0	0	0
			0	0	0	0
1 5/1	Brg/	1	0	0	0	0
			0	0	0	0

Group, Slot/Intf. These columns are described for the **vi** command on page 17-51.

Service/Instance. The Service Type (**Service**) and Instance (**Instance**) of this Service Type in the switch.

Service Type values are as follows (note that OmniStack switches do not provide FDDI or Frame Relay ports):

- **Rtr** Virtual router port
- **Brg** Virtual bridge port
- **Tnk** Virtual trunk port (used for ATM, FDDI, and WAN)
- **T10** 802.10 FDDI service port
- **FRT** Frame Relay trunk port
- **Lne** LAN Emulation service port
- **Vlc** VLAN clusters (X-LANE) service port
- **CIP** Classical IP service port

The Instance (**Inst**) is an identifier of this type of service within the switch. For example, if more than one virtual router port is configured in the switch, then each “instance” of a router will be given a different number.

Frames In/Out. The number of frames received or sent from this port. The top number for each port row is the number of frames received, and the bottom number is the number of frames sent. Statistics are not provided for virtual router ports in this display, but they are provided through Networking menu commands. See Chapters 23 and 25 for further information on router port statistics.

Octets In/Out. The number of octets, or bytes, received or sent from this port. The top number for each port row is the number of octets received, and the bottom number is the number of octets sent. Statistics are not provided for virtual router ports, but they are provided through Networking menu commands. See Chapters 23 and 25 for further information on router port statistics.

Ucast Pkts In/Out. The total number of unicast packets received or sent from this port. The top number for each port row is the number of unicast packets received, and the bottom number is the number of unicast packets sent. Statistics are not provided for virtual router ports, but they are provided through Networking menu commands. See Chapters 23 and 25 for further information on router port statistics.

Non Ucast Pkts In/Out. The total number of non-unicast packets received or sent from this port. Non-unicast frames include multicast and broadcast frames. The top number for each port row is the number of non-unicast packets received, and the bottom number is the number of non-unicast packets sent. Statistics are not provided for virtual router ports, but they are provided through Networking menu commands. See Chapters 23 and 25 for further information on router port statistics.

Viewing Port Errors

The **ve** command displays port error statistics for ports in the switch. Entering

```
ve
```

displays error statistics for all virtual ports in the switch. You can also display errors statistics for only ports in a specific Group by specifying the Group ID after the **ve** command. For example, to view errors only for ports in Group 6, you would enter

```
ve 6
```

You can also display error statistics for a specific port by entering the slot and port number after the **ve** command. For example, to view errors only for Port 1 on the module in Slot 4, you would enter

```
ve 4/1
```

The same type of information is displayed for a single Group or port as is displayed for all ports in a switch. The following screen shows a sample from the **ve** command when specified without any Group or port parameters.

Virtual Interface Error Information- For All Interfaces

Group	Slot/ Intf	Service/ Instance	Buffer Discards In	Out	Error Discards In	Out
2	All	Rtr/ 1				
3	All	Rtr/ 2				
1	All	Rtr/ 1				
1	3/1	Tnk/ 1	0	0	0	0
1	4/1	Brg/ 1	0	0	0	0
1	4/2	Brg/ 1	0	0	0	0
1	4/3	Brg/ 1	0	0	0	0
1	4/4	Brg/ 1	0	0	0	0
1	4/5	Brg/ 1	0	0	0	0
1	4/6	Brg/ 1	0	0	0	0
1	4/7	Brg/ 1	0	0	0	0
1	4/8	Brg/ 1	0	0	0	0
1	5/1	Brg/ 1	0	0	0	0

Group, Slot/Intf. These columns are described for the **vi** command on page 17-51.

Service/Instance. The Service Type (**Service**) and Instance (**Instance**) of this Service Type in the switch. Service Type values are as follows (note that OmniStack switches do not provide FDDI or Frame Relay ports):

- **Rtr** Virtual router port
- **Brg** Virtual bridge port
- **Tnk** Virtual trunk port (used for ATM, FDDI, and WAN)
- **T10** 802.10 FDDI service port
- **FRT** Frame Relay trunk port
- **Lne** LAN Emulation service port
- **Vlc** VLAN clusters (X-LANE) service port
- **CIP** Classical IP service port

The Instance (**Inst**) is an identifier of this type of service within the switch. For example, if more than one virtual router port is configured in the switch, then each “instance” of a router will be given a different number.

Buffer Discards In/Out. For transmit (**Out**) and receive (**In**), the number of frames discarded due to a lack of buffer space. Buffer discard information is not provided for virtual router ports.

Error Discards In/Out. For transmit (**Out**) and receive (**In**), the number of frames discarded due to errors. Error discard information is not provided for virtual router ports.

Port Mirroring

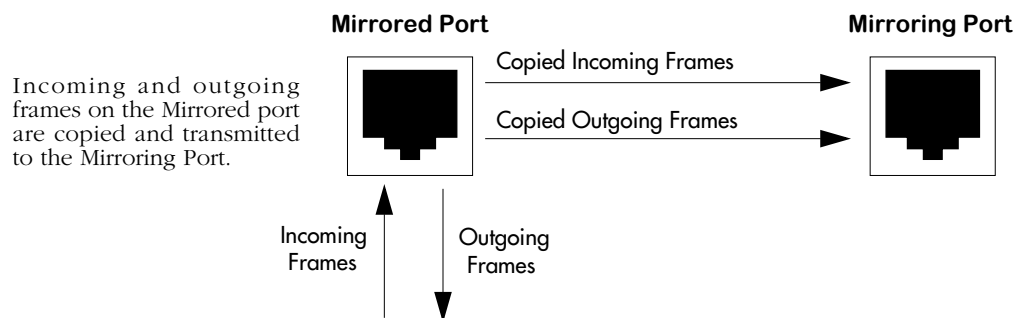
You can set up Port Mirroring for any pair of Ethernet (10 or 10/100 Mbps) ports within the same switch. When you enable port mirroring, the active, or “mirrored,” port transmits and receives network traffic normally, and the “mirroring” port receives a copy of all transmit and receive traffic to the active port. You can connect an RMON probe or network analysis device to the mirroring port to see an exact duplication of traffic on the mirrored port without disrupting network traffic to and from the mirrored port.

Port mirroring is supported on OmniStack switches for Ethernet (10 or 10/100 Mbps) ports only. An Ethernet port can only be mirrored by one other Ethernet port. A mirroring port can only mirror one port at a time. Up to five (5) mirroring sessions (mirrored-mirroring port pairs) are supported in a single switch. The mirrored and mirroring ports can be in different Groups and different VLANs.

How Port Mirroring Works

When a frame is received on a Mirrored Port it is copied and sent to the Mirroring Port. The received frame is actually transmitted twice across the switch backplane—once for normal bridging and then again to the Mirroring Port.

When a frame is transmitted by the mirrored port, a copy of the frame is made, tagged with the mirroring port as the destination, and sent back over the switch backplane to the mirroring port. The following diagram illustrates the data flow for a Mirrored-Mirroring port pair.



Relationship Between Mirrored and Mirroring Port

When port mirroring is enabled, there may be some performance degradation since all frames received and transmitted by the Mirrored port need to be copied and sent to the Mirroring port.

What Happens to the Mirroring Port

Once you set up port mirroring and attach cables to the Mirrored and Mirroring ports, the Mirroring port is administratively disabled and no longer a part of the Bridging Spanning Tree. The Mirroring port does not transmit or receive any traffic on its own. In addition, the Admin Status of the mirroring port displays in switch software commands, such as **vi**, as

M <slot> <port>

where **<slot>** is the slot number of the module containing the mirrored port, and **<port>** is the port number of the mirrored port. For example, if the Admin Status of a port displayed as

M 3 02

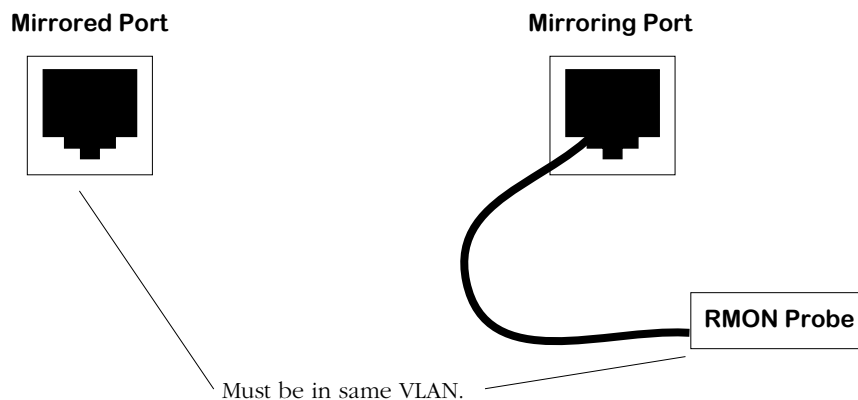
then you would know this port is mirroring traffic for Port 2 in Slot 3.

If a cable is not attached to the Mirrored port, port mirroring will not take place. In this case, the Mirroring Port reverts back to its normally operational state and will bridge frames as if port mirroring were disabled.

Using Port Mirroring With External RMON Probes

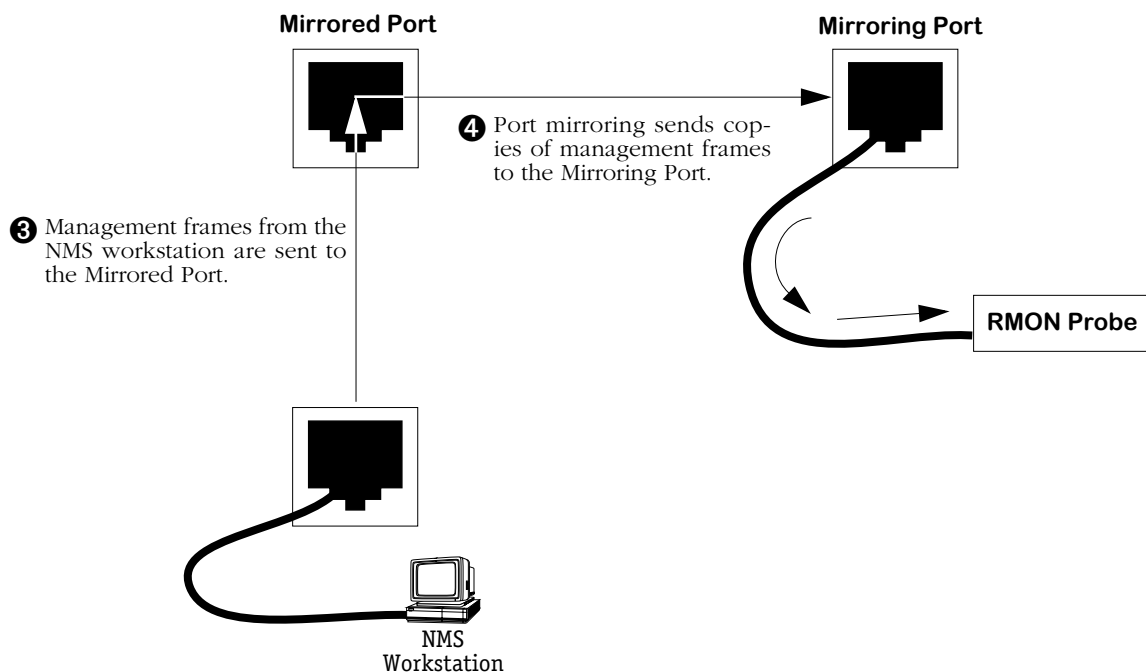
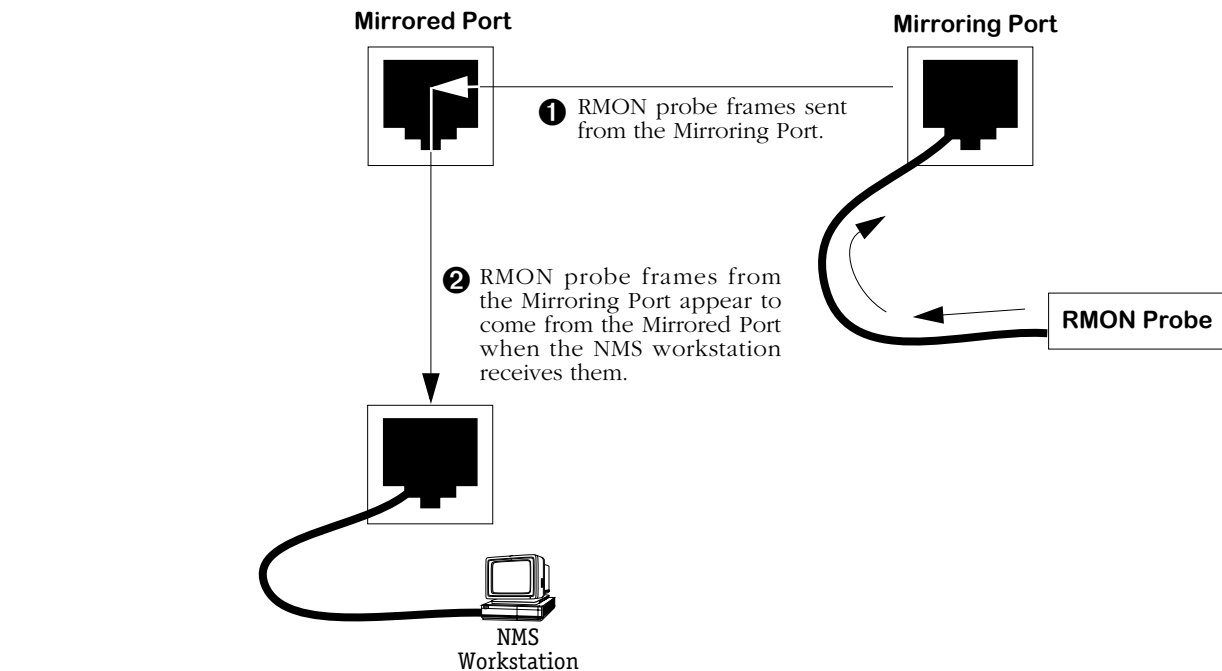
Port mirroring is a helpful monitoring tool when used in conjunction with an external RMON probe. Once you set up port mirroring, the probe can collect all relevant RMON statistics for traffic on the mirrored port. You can also move the Mirrored Port so that the Mirroring Port receives data from different ports. In this way, you can roam the switch and monitor traffic at various ports.

If you attach an external RMON probe to a mirroring port, that probe must have an IP address that places it in the same VLAN as the mirrored port. In addition if you change the mirrored port, then you must again make sure that the RMON probe is in the same VLAN as that new mirrored port.



Mirrored and Mirroring Ports in Same VLAN

You can view the Mirroring Port as a physical extension of the Mirrored Port. Frames received from an RMON probe attached to the Mirroring Port can be seen as being received by the Mirrored Port. These frames from the Mirroring Port are marked *as if they are received on the Mirrored Port* before being sent over the switch backplane to an NMS station. Therefore, management frames from an NMS station that are destined for the RMON probe are first forwarded out the Mirrored Port. After being received on the Mirrored Port, copies of the frames are mirrored out the Mirroring Port—the probe attached to the Mirroring Port receives the management frames. The following illustration shows this data flow.



Port Mirroring Using an External RMON Probe

Setting Up Port Mirroring

You set up port mirroring when you add or modify a port through the **addvp** or **modvp** commands. The switch software senses the type of port you are configuring, so it will only prompt you for port mirroring when configuring an Ethernet port. Follow the steps below to set up port mirroring.

1. Start the **addvp** or **modvp** command for the virtual port that you want to mirror.
2. At the **Command** prompt enter **8=e**, press **<Enter>** and you will be prompted for the slot and port number of the “mirroring” port (i.e., the port that can “see” all traffic for this port):

Mirroring vport slot/port ? () :

3. Enter the mirroring port’s slot, a slash (/), the port number, and then press **<Enter>**. The port that you indicate here will be disabled and only capable of receiving duplicate traffic from the mirrored port. If port mirroring is not supported on this port, then the following prompt will display:

mirroring not supported on this port type

After entering the Mirroring slot and port number, the **addvp** or **modvp** screen of options re-displays with the changes you entered. If you are done modifying or adding the port, enter **save** at the **Command** prompt. If using the **addvp** command a message indicating that you have successfully set up the port displays. Port mirroring takes place immediately, so you could now connect a probe or network analyzer to the Mirroring port.

Disabling Port Mirroring

You can disable port mirroring through the **modvp** command. Follow these steps to disable port mirroring.

1. Start the **modvp** command for the virtual port on which you want to disable port mirroring.
2. At the **Command** prompt enter **8=d**, press **<Enter>**. The **modvp** screen re-displays. The **Mirrored Port Status** field should read **Disabled, available**.

Port Monitoring

An essential tool of the network engineer is a network packet capture device. A packet capture device is usually a PC-based computer, such as the Sniffer[®], that provides a means for understanding and measuring data traffic of a network. Understanding data flow in a VLAN-based switch presents unique challenges primarily because traffic takes place *inside* the switch, especially on dedicated devices.

The port monitoring feature built into OmniStack software allows the network engineer to examine packets to and from a specific Ethernet port. Port monitoring has the following features:

- Software commands to enable and display captured port data.
- Captures data in Network General® file format.
- Limited protocol parsing (basic IP protocols and IPX) in console dump display.
- Data packets time stamped.
- One port monitored at a time.
- RAM-based file system.
- Memory buffer space from 1 MB to 8 MB.
- Statistics gathering and display
- Monitors only Ethernet ports
- Filtering limited to basic packet type—broadcast, multicast or unicast.

You can select to dump real-time packets to the terminal screen, or send captured data to a file. Once a file is captured, you can FTP it to a Sniffer for viewing.

Port Mirroring

An alternate method of monitoring ports is Port Mirroring, which allows a network engineer to attach a Sniffer to one Ethernet port and mirror traffic to and from any other Ethernet port. Port mirroring is described in *Port Mirroring* on page 17-58.

Port Monitoring Menu

The port monitoring commands are contained on the port monitoring menu, which is a sub-menu of the Networking menu. The port monitoring menu displays as follows:

Command	Port Monitoring Menu				
pmon	Port monitor utility				
pmcfg	Configure port monitor parameters				
pmstat	View port monitor statistics				
pmd	Port monitor disable				
pmp	Port monitor pause				
	Main Interface	File Security	Summary System	VLAN Services	Networking Help

The commands in this menu are described in the following sections.

RAM Disk System for Data Capture Files

Port monitoring uses a RAM disk for fast temporary storage of data capture files. The RAM disk has a separate directory designation of **/ram**. RAM-based files are created in DOS-FAT format and they are displayed in UPPERCASE.

You can copy files between the **/ram** disk system and the standard **/flash** file system. In addition, files in the RAM disk system are retrievable via FTP. Both the **/ram** file system and the **/flash** file system are accessible by using the UNIX/DOS-style change directory (**cd**) command.

Note

The RAM drive is part of DRAM memory. If you power off or reboot the switch, any files saved in the RAM drive will be lost.

Configuring RAM Drive Resources (pmcfg)

The **pmcfg** command allows you to select the size of the RAM disk file system or to delete the RAM disk. In addition, it allows you to configure the amount of data collected for each packet capture. To begin configuring RAM drive resources, enter

```
pmcfg
```

A screen similar to the following displays:

```
RAM disk size : 1000 Kilobytes
Lines displayed: 1
Change any of the above (y/n)? (n)
```

To change one of the settings, enter a **Y** and press **<enter>**. You will be prompted for a new RAM drive size. Select a size in kilobytes between 1000 and 8000. You can also delete the RAM drive by entering a size of zero (0). Changing the RAM disk size also requires that you reboot the system.

The **Lines displayed** controls the amount of data displayed to the terminal when you choose to dump session data to the computer screen. You can specify the number of lines to display while viewing port monitor data on the screen.

Changing the Default System Directory (cd)

After a port monitoring session is enabled the default directory is the RAM disk system (**/ram**). To switch back to the standard default flash file system (**/flash**) use the **cd** command. To switch back to the default directory, enter

```
cd /flash
```

To switch back to the RAM disk directory, enter

```
cd /ram
```

Starting a Port Monitoring Session (pmon)

You enable a port monitoring session through the **pmon** command. To start a session, enter **pmon** followed by the slot and port number that you want to monitor. For example, to monitor a port that is the first port in the fourth slot of the switch, you would enter

```
pmon 4/1
```

You can only monitor Ethernet ports. If a port is already being mirrored (enabled through the **addvp** or **modvp** command) you cannot monitor it. Also, you cannot set up more than one monitoring session on the same port.

If the port is currently being monitored, or mirrored, the following message displays:

```
Port 4/1 is being monitored.  
Disable monitoring? (y)
```

If the port is not being monitored, or mirrored, the following message displays:

```
Port 4/1 is not being monitored, or mirrored.  
Enable monitoring? (y)
```

Enter a **Y** and press **<enter>** at this prompt. The following screen of options displays:

```
Slot/Port           : 4/1  
RAM disk size       : 1000 Kilobytes  
Capture to filename : y  
Capture filename    : PMONITOR.ENC  
Dump to screen      : y  
Broadcast frames    : y  
Multicast frames    : y  
Unicast frames      : y  
Change any of the above (y/n)? (n) :
```

If you want to change any of the values, enter a **Y** and press **<enter>**. You will be prompted for all of the values in the screen except the **RAM disk size**, which you must change through the **pmcfg** command before starting the session. The information selected in this screen will be saved in flash configuration memory.

Enter any new values as prompted. The above screen re-displays to show the new values. Press **<enter>** to accept the updated values. Messages similar to the following display:

```
1048576 byte RAM drive /ram already initialized.  
Bytes remaining on RAM disk = 1040384
```

The port monitoring session has begun. What happens at this point depends on whether you chose the **Dump to screen** option. The sections below describe what happens in each case.

◆ Important Note ◆

If you change the capture filename from the default, you must specify **/ram**. Otherwise, the file will be saved in the flash directory.

If You Chose *Dump to Screen*

If you selected the **Dump to screen** option, then a real-time synopsis of the session displays on your terminal screen. The following shows an example of this data

```
Enter 'p' to pause, 'q' to quit.
Destination      | Source          | Type | Data
-----
00:20:DA:04:01:02 | 00:20:DA:04:01:01 | ICMP | 01:02:03:04:05:06:07:08
00:20:DA:04:01:02 | 00:20:DA:04:01:01 | ICMP | 01:02:03:04:05:06:07:08
FF:FF:FF:FF:FF:FF | 00:20:DA:02:10:E3 | ARP-C | 08:06:00:01:08:00:06:04
FF:FF:FF:FF:FF:FF | 00:20:DA:6F:97:A3 | RIP   | 08:00:45:00:00:34:22:30
```

Each line in the display represents a packet. The destination MAC address, source MAC address, protocol type and actual packet data are shown. The amount of data shown is configured through the **pmcfg** command. The above sample shows 16 bytes of data per packet. You can stop the data dump to the screen at anytime by pressing **q** to quit. You can also pause the data dump by pressing **p** to pause.

If You Did Not Choose *Dump to Screen*

If you did not select the **Dump to screen** option, then the system prompt will return and port monitoring occurs in the background. You can continue using other UI commands. The port monitoring session data is saved in the file you indicated through the **pmmon** screen. You can monitor the session at anytime by using the **pmstats** command. You can also end or pause an in-progress session using the **pmdelete** or **pmppause** commands, respectively. The following sections describes **pmdelete** and **pmppause**.

Ending a Port Monitoring Session (**pmdelete**)

The **pmdelete** command ends a port monitoring data capture session that is being saved to file but not being dumped to the console screen. To end the session, enter:

```
pmd
```

A message similar to the following displays:

```
Port monitoring session terminated, data file is xxxxx.ENC.
```

If a port monitoring session was not in progress then the following message displays:

```
No ports being monitored.
```

Pausing a Port Monitoring Session (**pmppause**)

The **pmppause** command pauses a port monitoring data capture session that is being saved to file but not being dumped to the console screen. To pause the session, enter:

```
pmpp
```

The following message displays

```
Pausing monitor data capture/display.
```

To resume the port monitoring session, enter **pmpp** again. The following message displays:

```
Resuming monitor data capture.
```

If a port monitoring session was not in progress, then the following message would display:

```
No ports being monitored.
```

Ending a Port Monitoring Session

After you quit a port monitoring session, the default directory changes to **/ram** and the current files on the RAM drive are listed. The screen below shows an example of the display at the completion of a monitoring session.

**Port monitoring capture done. Current capture files listed:
Current working directory '/ram'.**

**PM0302.ENC 65536 10/20/96 12:12
PM0303.ENC 32768 10/20/96 11:15**

950272 bytes free

Viewing Port Monitoring Statistics (pmstat)

The **pmstat** command displays the statistics gathered for the current or most recent port monitoring session. If a port monitoring session is currently in progress, then it displays the results of the in-progress session. If a port monitoring session is not in progress, then it displays results of the most recently completed session. To view session statistics, enter

pmstat

A screen similar to the following displays:

**Viewing capture statistics:
Percent RAM available: 96%
Frame type #Frames

Broadcast 108
Multicast 253
Unicast 301**

The **Percent RAM available** indicates how much of the configured RAM disk has been used by this port monitoring session. You can configure the size of the RAM disk through the **pmcfg** command; the default size is 1 MB. The remaining items in the display show the number of packets passed on the port broken down into broadcast, multicast, and unicast frames.

Port Mapping

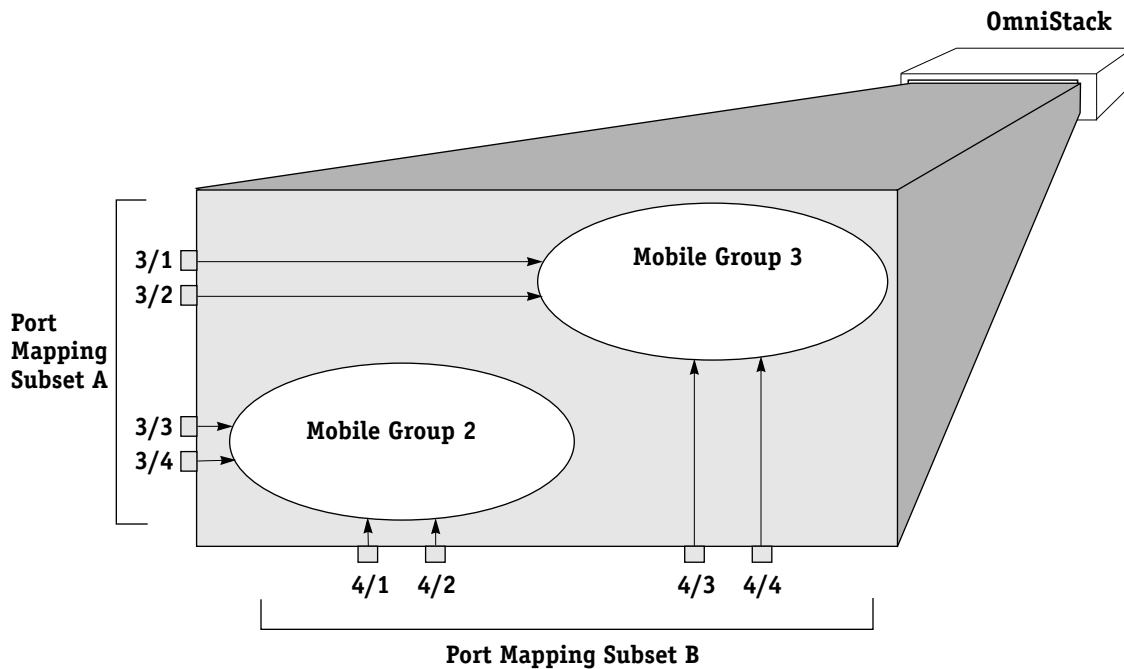
The OmniStack began as an any-to-any switching device, connecting different LAN interfaces, such as Ethernet, Token Ring, and FDDI. As networks grew and the traffic on them increased, a need arose for controlling some traffic, such as broadcasts. Virtual LANs, or VLANs, were introduced to segment traffic such that devices could only engage in switched communication with other devices in the same VLAN.

Some applications today require a further degree of traffic segmentation than that provided by VLANs. The port mapping feature allows you to further segment traffic *within* a VLAN or group by isolating a set of ports.

Groups/VLANs and Port Mapping

Port mapping does *not* affect existing group or AutoTracker VLAN operations in a switch. Group and VLAN membership are checked and applied before port mapping constraints are applied. Therefore, any constraints applied by port mapping only limit traffic flow *within* a group or VLAN; port mapping parameters do not provide any additional connectivity to a port. So if you add a port to a port mapping set, that port will be first subject to the constraints of its Group/VLAN and then the restrictions imposed by port mapping. Up to 128 port mapping sets can be configured per switch.

The illustration below helps show how group and port mapping constraints interact. The ports 3/1, 3/2, 4/3, and 4/4 are part of groups 3. By group membership, all of these ports have switched communication with each other. Likewise, the ports 3/3, 3/4, 4/1, and 4/2 have switched communication with each other as they all belong to group 2.



Groups and Port Mapping

Once a port mapping set is constructed, communication within each of the groups becomes more restricted. A port mapping set consists of *ingress* and *egress* ports; ingress ports can only send traffic to egress ports. In the above figure, all ports on slot 3 are ingress ports and ports on slots 4 are egress ports.

Port communication is uni-directional. A mapping between an ingress port and an egress port can only pass data from the ingress port to the egress port. To allow traffic to flow the from the egress port to the ingress port, it is necessary to create a new mapping.

This configuration restricts each port to communication *only with the other four ports in the opposite port mapping subset within the same group*. For example, port 3/1 can only send traffic to ports 4/3 and 4/4. It can no longer communicate with port 3/2 even though it is part of the same group. Port mapping restricts ports from communicating with other ports within the same subset.

Port mapping does not affect other ports in the group that are not part of the port mapping set.

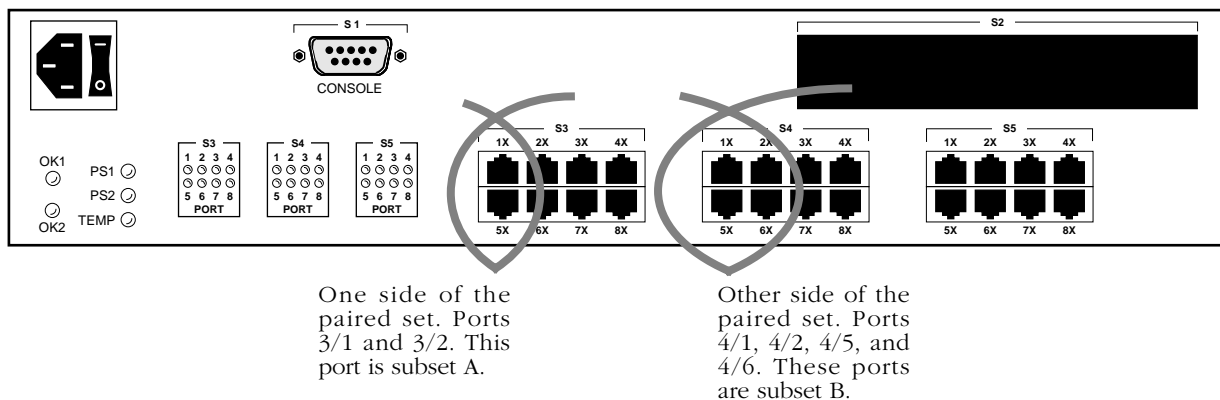
The Details of Port Mapping

Port mapping can be thought of as special rule that is applied after standard group and VLAN rules are applied. This rule statically assigns a port as either an ingress or egress port. Ingress ports can only communicate with egress ports. In this sense, one subset of ports is “mapped” to another subset of ports. Ports within the same subset can not communicate with each other or with another switch port that is not a member of the opposite port mapping subset.

Note

Port mapping restrictions are only applied to ports on 10/100 Ethernet ports.

As an illustration, see the diagram of an OmniStack. The ports are in slots 3 and 4. The ports that are circled are included in a port mapping subset. The first subset is port 3/1 and 3/5, and are ingress ports. The second subset includes ports 4/1, 4/2, 4/5 and 4/6, and are egress ports in the port mapping set.



Port Subsets in the Port Mapping Set

Who Can Talk to Whom?

The following matrix outlines which ports can communicate with each other in the example shown on the previous page *assuming all ports are part of the same group or VLAN*. A port can only communicate with ports in the opposite subset within the port mapping set.

Switch Ports That May Communicate

	3/1	3/2	4/1	4/2	4/5	4/6
3/1	N/A	No	Yes	Yes	Yes	Yes
3/2	No	N/A	Yes	Yes	Yes	Yes
4/1	Yes	Yes	N/A	No	No	No
4/2	Yes	Yes	No	N/A	No	No
4/5	Yes	Yes	No	Yes	N/A	No
4/6	Yes	Yes	No	No	No	N/A

Port communication is uni-directional. A mapping between an ingress port and an egress port can only pass data from the ingress port to the egress port. To allow traffic to flow the from the egress port to the ingress port, it is necessary to create a new mapping.

It's important to remember that the port mapping configuration is affected by existing group/VLAN rules. If the ports in the above example belonged to three groups based on IP network rules, then they would be restricted by group membership and port mapping.

Port Mapping Limitations

The following are restrictions to the use of the port mapping feature:

- Port mapping cannot be used with ports assigned to an 802.1Q group.
- Port mapping cannot be used with an OmniChannel unless all ports in the OmniChannel are included in the port mapping (on either the ingress or egress list). For example, if ports 3/1-3/4 are an OmniChannel, all four ports must be in the ingress or egress list. You could not just map port 3/1.

Creating a Port Mapping Set

Use the **pmapcr** command to create a port mapping set. Follow these steps:

1. Enter **pmapcr** at a system prompt.
2. The following screen displays:

Port Map Configuration

- 1. Ingress List :
- 2. Egress List :

Enter the ingress ports and egress ports for this map set. This is done by entering the line number, an equal sign, and the port (or ports) to be added. For example, if you want to create a map set with an ingress port of 3/6 and an egress port of 3/8, you would enter the following at the prompt:

```
1=3/6
2=3/8
```

This must be done in two separate operations, one for the ingress and one for the egress lists. You can add more than one port to a list by using a comma (,) between slot/port designations, or a dash (-) between port numbers. For example, if you wanted to make ports 3/1, 3/6, 3/7, 3/8, and 3/9 egress ports for this map set, you would enter the following:

```
2=3/1, 3/6-9
```

A switch port in the ingress list can only communicate with switch ports in the egress list. Switch ports in the same list cannot communicate with each other or any other ports in the switch. For example, if you enter:

```
1=3/1, 3/2
2=3/3, 3/4
```

then you are creating a paired set of four ports. Port 3/1 can only communicate with ports 3/3 or 3/4. It cannot communicate with any other ports in the switch, including port 3/2. Port 3/1 also can only communicate with ports 3/3 and 3/4, but no others.

Any port type may be added to a port mapping set. However, only Mammoth-generation Ethernet ports will be restricted by port mapping limitations. For example, you could add a non-Ethernet port to the set, but traffic from that port would not be restricted.

3. You will want to save your configuration, so enter an **s** at the **port-mapping** prompt. Your configuration will be saved. A prompt similar to the following appears to confirm the creation of the port map:

Port Map 7 created.

The port map number is used when modifying the map set.

It is important to remember that port communication is uni-directional. A mapping between an ingress port and an egress port can only pass data from the ingress port to the egress port. To allow traffic to flow from the egress port to the ingress port, it is necessary to create a new mapping.

Adding Ports to a Port Mapping Set

You can add ports to a port map set once it has been created using the **pmapmod** command. Follow these steps:

1. Enter the **pmapmod** command at a system prompt, as shown:

```
pmapmod <pmap id>
```

where **<pmap id>** is the map set number shown when the map set was created. (To view a list of all existing map sets, see *Viewing a Port Mapping Set* on page 17-73.) For example, to modify map set 5, you would enter the following:

```
pmapmod 5
```

2. The following screen displays:

```

Port Mapping Configuration
=====
Port Map Id      Ingress Ports      Egress Ports
-----
          5      3/1, 3/2, 3/3      3/1, 3/2, 3/3

Modify Port Map 5

1. Add Ports to Ingress List      :
2. Add Ports to Egress List      :
3. Delete Ports from Ingress List :
4. Delete Ports from Egress List  :
5. View Port Map Configuration   :
```

Note that the current ports in the port mapping set are displayed. Use this information to make decisions on the ports you want to add or remove from the set.

Enter the line number for the operation you want to perform (a **1** for the ingress list or a **2** for the egress list), an equal sign (=), and the ports to be added. For example, add port 3/2 to the ingress list and the egress list, enter the following (in two separate operations):

```
1=3/2
2=3/2
```

You can add more than one port to a list by using a comma (,) between slot/port designations, or a dash (-) between port numbers. For example, if you wanted to make ports 4/1, 4/6, 4/7, 4/8, and 4/9 egress ports for this map set, you would enter the following:

```
2=4/1, 4/6-9
```

3. To view the changes, enter a **5 (View Port Map Configuration)**, and equal sign (=), and a **y**, as shown:

```
5=y
```

This will refresh the Port Mapping Configuration screen and display any changes you have made.

4. Quit the session by entering a **q** at the prompt.

Removing Ports from a Port Mapping Set

You can remove ports to a port map set once it has been created using the **pmapmod** command. Follow these steps:

1. Enter the **modpmap** command at a system prompt, as shown:

```
pmapmod <pmap id>
```

where **<pmap id>** is the map set number shown when the map set was created. (To view a list of all existing map sets, see *Viewing a Port Mapping Set* on page 17-73.) For example, to modify map set 5, you would enter the following:

```
pmapmod 5
```

2. The Port Mapping Configuration screen displays (as shown above in *Adding Ports to a Port Mapping Set* on page 17-71).

Enter the line number for the operation you want to perform (a **3** for the ingress list or a **4** for the egress list), an equal sign (=), and the ports to be added. For example, remove port 3/2 to the ingress list and the egress list, enter the following (in two separate operations):

```
3=3/2  
4=3/2
```

You can remove more than one port to a list by using a comma (,) between slot/port designations, or a dash (-) between port numbers. For example, if you wanted to remove ports 4/1, 4/6, 4/7, 4/8, and 4/9 from the egress list of this map set, you would enter the following:

```
4=4/1, 4/6-9
```

3. To view the changes, enter a **5** (view port map configuration), and equal sign (=), and a **y**, as shown:

```
5=y
```

This will refresh the Port Mapping Configuration screen and display any changes you have made.

4. Quit the session by entering a **q** at the prompt.

Viewing a Port Mapping Set

You can view a port mapping set using the **vpmap** command. Enter the **pmapv** command as shown:

```
pmapv <pmap id>
```

where **<pmap id>** is the map set number shown when the map set was created. For example, to modify map set 5, you would enter the following:

```
pmapv 5
```

The following screen is shown:

```

Port Mapping Configuration
=====
Port Map Id      Ingress Ports      Egress Ports
-----
5                3/1, 3/2, 3/3     3/1, 3/2, 3/3

```

As a variation of this command, enter the **vpmap** command with no port map identification. This will display all port mapping sets configured for this switch.

Port Map Id. An identification number for the port map set, generated when the set is created.

Ingress Ports. The switch ports designated as ingress ports for this port map set. Ingress ports can only communicate with egress ports.

Egress Ports. The switch ports designated as egress ports for this port map set. Egress ports can only communicate with ingress ports.

Deleting a Port Mapping Set

You can delete a port mapping set after it is created. Enter **pmapdel** at a prompt as shown:

```
pmapdel <pmap id>
```

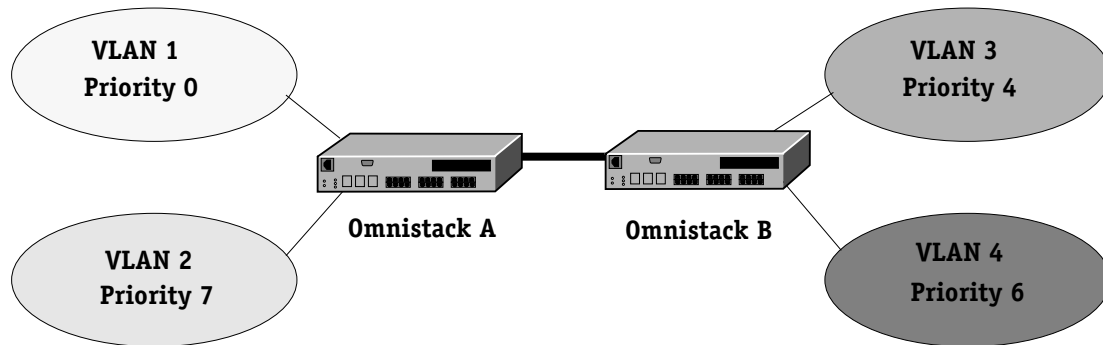
where **<pmap id>** is the map set number shown when the map set was created. (To view a list of all existing map sets, see *Viewing a Port Mapping Set* on page 17-73.) For example, to modify map set 5, you would enter the following:

```
pmapdel 5
```

Priority VLANs

Prioritizing VLANs allows you to set a value for traffic based on the destination VLAN of packets. Traffic with the higher priority destination will be delivered first. VLAN priority can be set from 0 to 7, with 7 being the level with the most priority.

The following diagram illustrates this idea:



In the above diagram, traffic from VLAN 1 to VLAN 4 would have priority over traffic from VLAN 1 to VLAN 3. Conversely, traffic sent from VLAN 4 to VLAN 2 would have priority over traffic from VLAN 4 to VLAN 1.

Group priority can be set when creating a group using the **crgrp** command. For more information on the **crgrp** command, see Chapter 17, “Managing Groups and Ports.”

Group priority can be modified or viewed using the **prty_mod** and **prty_disp** commands, detailed below.

◆ Note ◆

Although the range of VLAN priority is 0-7, the current implementation only supports two levels of priority. In other words, 0-3 is one level and 4-7 is another. Future releases will expand the number of priority levels.

Configuring VLAN Priority

To configure the priority of a VLAN:

1. Enter the **prty_mod** command at the system prompt, as shown:

```
prty_mod <groupId>
```

where **<groupId>** is the group number associated with the VLAN whose priority is being set. For example, to modify the priority of the VLAN for Group 2, you would enter the following:

```
prty_mod 2
```

The following prompt is shown:

```
Enter a priority value which is between 0 and 7: 0
```

2. Enter the number value that is to be the new priority level for this VLAN. The highest (most important) value is 7.
3. Press **<enter>**. A message similar to the following is displayed:

```
Priority for VLAN 2 has been set as 7
```

Viewing VLAN Priority

The priority level for all configured VLANs can be viewed by using the **prty_disp** command. Enter the **prty_disp** at the system prompt, as shown:

```
prty_disp <groupId>
```

where **<groupId>** is the group number associated with the VLAN whose priority is being viewed. For example, to view the priority of the VLAN for Group 2, you would enter the following:

```
prty_disp 2
```

A display similar to the following is shown:

```
The priority of group 2 is 7
```

As a variation of this command, you can enter **prty_disp** at the system prompt without a group number. This will display the priority of all VLANs.

