Configuring VLANs

This chapter provides the following information about configuring and monitoring 802.1Q VLANs on Matrix N-Series, Matrix X, SecureStack, and secure switch devices.

Note: This document describes the configuration and operation of VLANs as defined by the IEEE 802.1Q standard and assumes that all devices being configured support that standard. No other types of VLANs will be covered.

What Is a VLAN?

A VLAN is a Virtual Local Area Network — a grouping of network devices that is logically segmented by functions, project teams, or applications without regard to the physical location of users. For example, several end stations might be grouped as a department, such as Engineering or Finance, having the same attributes as a LAN, even though they are not all on the same physical LAN segment.

To accomplish this logical grouping, the network administrator uses 802.1Q VLAN-capable switching devices and assigns each switch port in a particular group to a VLAN. Ports in a VLAN share broadcast traffic and belong to the same broadcast domain. Broadcast traffic in one VLAN is not transmitted outside that VLAN.

Why Would I Use VLANs in My Network?

Virtual LANs allow you to partition network traffic into logical groups and control the flow of that traffic through the network. Once the traffic and, in effect, the users creating the traffic, are assigned to a VLAN, then broadcast and multicast traffic is contained within the VLAN and users can be allowed or denied access to any of the network's resources. Also, you have the option of configuring some or all of the ports on a device to allow frames received with a particular VLAN ID and protocol to be transmitted on a limited number of ports. This keeps the traffic associated with a particular VLAN and protocol isolated from the other parts of the network.

The primary benefit of 802.1Q VLAN technology is that it allows you to localize and segregrate traffic, improving your administrative efficiency, and enhancing your network security and performance.

Figure 1 shows a simple example of using port-based VLANs to achieve these benefits. In this example, two buildings house the Sales and Finance departments of a single company, and each building has its own internal network. The end stations in each building connect to a switch on the bottom floor. The two switches are connected to one another with a high speed link.

Figure 1 VLAN Business Scenario

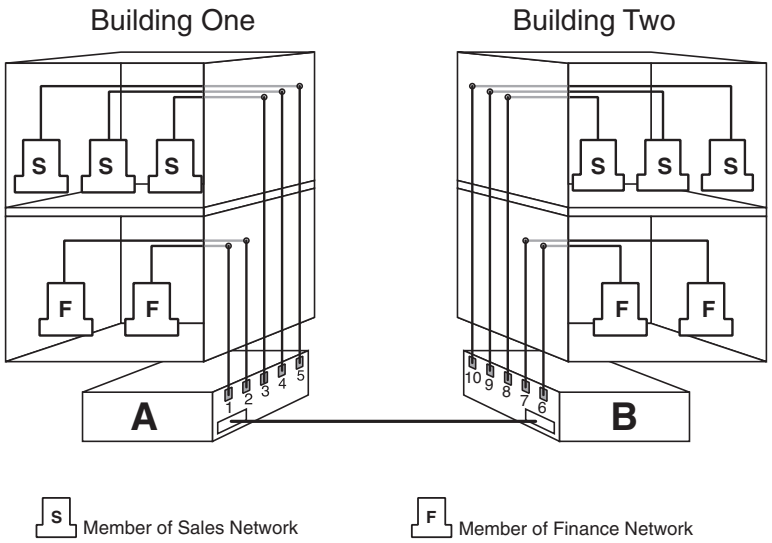

Without any VLANs configured, the entire network in the example in Figure 1 would be a broadcast domain, and the switches would follow the IEEE 802.1D bridging specification to send data between stations. A broadcast or multicast transmission from a Sales workstation in Building One would propagate to all the switch ports on Switch A, cross the high speed link to Switch B, and then be propagated out all switch ports on Switch B. The switches treat each port as being equivalent to any other port, and have no understanding of the departmental memberships of each workstation.

Once Sales and Finance are placed on two separate VLANs, each switch understands that certain individual ports or frames are members of separate workgroups. In this environment, a broadcast or multicast data transmission from one of the Sales stations in Building One would reach Switch A, be sent to the ports connected to other local members of the Sales VLAN, cross the high speed link to Switch B, and then be sent to any other ports and workstations on Switch B that are members of the Sales VLAN.

Another benefit to VLAN use in the preceding example would be your ability to leverage existing investments in time and equipment during company reorganization. If, for instance, the Finance users change location but remain in the same VLAN connected to the same switch port, their network addresses do not change, and switch and router configuration is left intact.

How Do I Implement VLANs?

By default, all Enterasys switches run in 802.1Q VLAN operational mode. All ports on all Enterasys switches are assigned to a default VLAN (VLAN ID 1), which is enabled to operate and assigns all ports an egress status of untagged. This means that all ports will be allowed to transmit frames from the switch without a VLAN tag in their header. Also, there are no forbidden ports (prevented from transmitting frames) configured.

You can use the CLI commands described in this document to create additional VLANs, to customize VLANs to support your organizational requirements, and to monitor VLAN configuration.

Preparing for VLAN Configuration

A little forethought and planning is essential to a successful VLAN implementation. Before attempting to configure a single device for VLAN operation, consider the following:

- What is the purpose of my VLAN design? (For example: security or traffic broadcast containment).
- How many VLANs will be required?
- What stations (end users, servers, etc.) will belong to them?
- What ports on the switch are connected to those stations?
- What ports will be configured as GARP VLAN Registration Protocol (GVRP) aware ports?

Determining how you want information to flow and how your network resources can be best used to accomplish this will help you customize the tasks described in this document to suit your needs and infrastructure.

Once your planning is complete, you would proceed through the steps described in "Configuring VLANs" on page 9.

Understanding How VLANs Operate

802.1Q VLAN operation differs slightly from how a switched networking system operates. These differences are due to the importance of keeping track of each frame and its VLAN association as it passes from switch to switch, or from port to port within a switch.

VLAN-enabled switches act on how frames are classified into a particular VLAN. Sometimes, VLAN classification is based on tags in the headers of data frames. These VLAN tags are added to data frames by the switch as the frames are transmitted out certain ports, and are later used to make forwarding decisions by the switch and other VLAN aware switches. In the absence of a VLAN tag header, the classification of a frame into a particular VLAN depends upon the configuration of the switch port that received the frame.

The following basic concepts of VLAN operation will be discussed in this section:

- Learning Modes and Filtering Databases (page 3)
- VLAN Assignment and Forwarding (page 4)
- Example of a VLAN Switch in Operation (page 5)

Learning Modes and Filtering Databases

Addressing information the switch learns about a VLAN is stored in the filtering database assigned to that VLAN. This database contains source addresses, their source ports, and VLAN IDs, and is referred to when a switch makes a decision as to where to forward a VLAN tagged frame. Each filtering database is assigned a Filtering Database ID (FID).

A switch learns and uses VLAN addressing information by the following modes:

- **Independent Virtual Local Area Network (VLAN) Learning (IVL)**: Each VLAN uses its own filtering database. Transparent source address learning performed as a result of incoming VLAN traffic is not made available to any other VLAN for forwarding purposes. This setting is useful for handling devices (such as servers) with NICs that share a common MAC address. One FID is assigned per VLAN. This is the default mode on Enterasys switches.
- **Shared Virtual Local Area Network (VLAN) Learning (SVL)**: Two or more VLANs are grouped to share common source address information. This setting is useful for configuring more complex VLAN traffic patterns, without forcing the switch to flood the unicast traffic in each direction. This allows VLANs to share addressing information. It enables ports or switches in different VLANs to communicate with each other (when their individual ports are configured to allow this to occur). One FID is used by two or more VLANs.

VLAN Assignment and Forwarding

Receiving Frames from VLAN Ports

By default, Enterasys switches run in 802.1Q operational mode, which means that every frame received by the switch must belong to, or be assigned to, a VLAN. The type of frame under consideration and the filter setting of the switch determines how it forwards VLAN frames. This involves processing traffic as it enters (ingresses) and exits (egresses) the VLAN switch ports as described below.

Untagged Frames

When, for example, the switch receives a frame from Port 1 and determines the frame does not currently have a VLAN tag, but recognizes that Port 1 is a member of VLAN A, it will classify the frame to VLAN A. In this fashion, all untagged frames entering a VLAN switch assume membership in a VLAN.

Note: A VLAN ID is always assigned to a port. By default, it is the default VLAN (VLAN ID = 1).

The switch will now decide what to do with the frame, as described in "Forwarding Decisions" on page 5.

Tagged Frames

When, for example, the switch receives a tagged frame from Port 4 and determines the frame is tagged for VLAN C, it will classify it to that VLAN regardless of its port VLAN ID (PVID). This frame may have already been through a VLAN aware switch, or originated from a station capable of specifying a VLAN membership. If a switch receives a frame containing a tag, the switch will classify the frame in regard to its tag rather than the PVID for its port, following the ingress precedence rules listed below.

Ingress Precedence

VLAN assignment for received (ingress) frames is determined by the following precedence:

1. 802.1Q VLAN tag (tagged frames only)
2. Policy or Traffic Classification (which may overwrite the 802.1Q VLAN tag) For more information, refer to "Configuring Protocol-Based VLAN Classification" on page 15.
3. Port VID (PVID)

Forwarding Decisions

VLAN forwarding decisions for transmitting frames is determined by whether or not the traffic being classified is or is not in the VLAN's forwarding database as follows:

- **Unlearned traffic**: When a frame's destination MAC address is not in the VLAN's forwarding database (FDB), it will be forwarded out of every port on the VLAN's egress list with the frame format that is specified. Refer to"Broadcasts, Multicasts, and Unlearned Unicasts"below for an example.
- **Learned traffic**: When a frame's destination MAC address is in the VLAN's forwarding database, it will be forwarded out of the learned port with the frame format that is specified. Refer to "Learned Unicasts" below for an example.

Broadcasts, Multicasts, and Unlearned Unicasts

If a frame with a broadcast, multicast, or other unknown address is received by an 802.1Q VLAN aware switch, the switch checks the VLAN classification of the frame. The switch then forwards the frame out all ports that are identified in the Forwarding List for that VLAN. For example, if Port 3, shown in the example in Figure 2, received the frame, the frame would then be sent to all ports that had VLAN C in their Port VLAN List.

Learned Unicasts

When a VLAN switch receives a frame with a known MAC address as its destination address, the action taken by the switch to determine how the frame is transmitted depends on the VLAN, the VLAN associated FID, and if the port identified to send the frame is enabled to do so.

When a frame is received it is classified into a VLAN. The destination address is looked up in the FID associated with the VLAN. If a match is found, it is forwarded out the port identified in the lookup if, and only if, that port is allowed to transmit frames for that VLAN. If a match is not found, then the frame is flooded out all ports that are allowed to transmit frames belonging to that VLAN.

Example of a VLAN Switch in Operation

The operation of an 802.1Q VLAN switch is best understood from a point of view of the switch itself. To illustrate this concept, the examples that follow view the switch operations from *inside* the switch.

Figure 2 depicts the inside of a switch with six ports, numbered 1 through 6. The switch has been configured to associate VLAN A and B with FID 2, VLAN C and D with FID 3, and VLAN E with FID 4. It shows how a forwarding decision is made by comparing a frame's destination MAC to the FID to which it is classified.

Figure 2 Inside the Switch

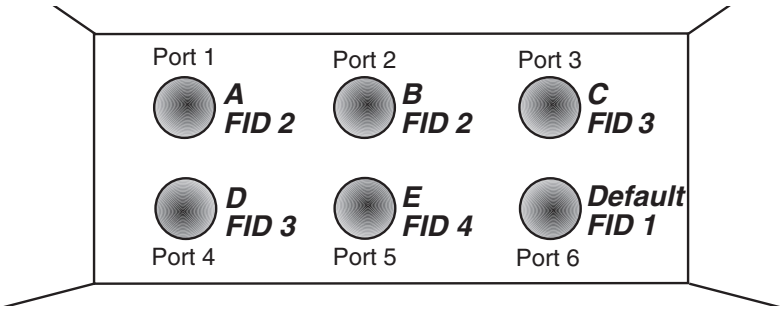

Assume a unicast untagged frame is received on Port 3 in the example in Figure 2. The frame is classified for VLAN C (the frame's PVID is VLAN C). The switch would make its forwarding decision by comparing the destination MAC address to information previously learned and entered into its filtering database. In this case, the MAC address is looked up in the FDB for FID 3, which is associated with VLANs C and D. Let's say the switch recognizes the destination MAC of the frame as being located out Port 4.

Having made the forwarding decision based on entries in the FID, the switch now examines the port VLAN egress list of Port 4 to determine if it is allowed to transmit frames belonging to VLAN C. If so, the frame is transmitted out Port 4. If Port 4 has not been configured to transmit frames belonging to VLAN C, the frame is discarded.

If, on the other hand, a unicast untagged frame is received on Port 5, it would be classified for VLAN E. Port 5 has is own filtering database and is not aware of what addressing information has been learned by other VLANs. Port 5 looks up the destination MAC address in its FID. If it finds a match, it forwards the frame out the appropriate port, if and only if, that port is allowed to transmit frames for VLAN E. If a match is not found, the frame is flooded out all ports that are allowed to transmit VLAN E frames.

VLAN Support on Enterasys Switches

Depending on the product family, Enterasys switches support a maximum of up to 4094 active VLANs. There is a distinction, however, between the maximum number of active VLANs some switches support and the range of VLAN ID (VID) values. For example, while the SecureStack and secure switch products support 1024 active VLANs, they do support VIDs from anywhere in the full 802.1Q specified range. These differences are listed below.

Maximum Active VLANs

The total number of active VLANs supported on Enterasys switch product families is:

- **Up to 4094** on Matrix N-Series and Matrix X
- **Up to 1024** on SecureStack and secure switch devices

Configurable Range

The allowable user-configurable range for VLAN IDs (VIDs) on Enterasys switches is from 2 through 4094. This range is based on the following rules:

- **VID 0** is the null VLAN ID, indicating that the tag header in the frame contains priority information rather than a VLAN identifier. It cannot be configured as a port VLAN ID (PVID).
- **VID 1** is designated the default PVID value for classifying frames on ingress through a switched port. This default can be changed on a per-port basis.
- **VID 4095** is reserved by IEEE for implementation use.

Notes: Each VLAN ID in a network must be unique. If you enter a duplicate VLAN ID, the Enterasys switch assumes you intend to modify the existing VLAN.

VLAN Types

Enterasys switches support traffic classification for the following VLAN types:

Static and Dynamic VLANs

All VLANs on an Enterasys switch are categorized as being either static or dynamic. Static VLANs are those that are explicitly created on the switch itself, persistently remaining as part of the configuration, regardless of actual usage. Dynamic VLANs, on the other hand, are not necessarily persistent. Their presence relies on the implementation of GVRP and its effect on egress membership as described in "GARP VLAN Registration Protocol (GVRP) Support" on page 7.

Port-Based VLANs

Port-based VLANs are configured by associating switch ports to VLANs in two ways: first, by manipulating the port VLAN ID (PVID); and second, by adding the port itself to the egress list of the VLAN corresponding to the PVID. Any traffic received by a port is associated to the VLAN identified by the port's PVID. By virtue of this association, this traffic may egress the switch only on those ports listed on the VLAN's egress list. For example, given a VLAN named "Marketing," with an ID value of 6, by changing the PVID values of ports 1 through 3 to 6, and adding those ports to the egress list of the VLAN, we effectively restrict the broadcast domain of Marketing to those three ports. If a broadcast frame is received on port 1, it will be transmitted out ports 2 and 3 only. In this sense, VLAN membership is determined by the location of traffic ingress, and from the perspective of the access layer—where users are most commonly located—egress is generally untagged.

Policy-Based VLANs

Rather than making VLAN membership decisions simply based on port configuration, each incoming frame can be examined by the classification engine which uses a match-based logic to assign the frame to a desired VLAN. For example, you could set up a policy which designates all e-mail traffic between the management officers of a company to a specific VLAN so that this traffic is restricted to certain portions of the network. With respect to network usage, the administrative advantages of policy classification would be application provisioning, acceptable use policy, and distribution layer policy. All of these provisions may involve simultaneous utilization of inter-switch links by multiple VLANs, requiring particular attention to tagged, forbidden, and untagged egress settings.

As described above, PVID determines the VLAN to which all untagged frames received on associated ports will be classified. Policy classification to a VLAN takes precedence over PVID assignment if:

- policy classification is configured to a VLAN, and
- PVID override has been enabled for a policy profile, and assigned to port(s) associated with the PVID.

For more information, refer to the Policy Classification chapter in your device's configuration guide.

GARP VLAN Registration Protocol (GVRP) Support

The purpose of the GARP (Generic Attribute Registration Protocol) VLAN Registration Protocol (GVRP) is to dynamically create VLANs across a switched network. GVRP allows GVRP-aware devices to dynamically establish and update their knowledge of the set of VLANs that currently have active members.

By default, GVRP is globally enabled on Enterasys switches and disabled at the port level. To allow GVRP to dynamically create VLANs, it must be enabled globally as well as on each individual port as described in "Configuring Dynamic VLANs" on page 15.

How It Works

When a VLAN is declared, the information is transmitted out GVRP configured ports on the device in a GARP formatted frame using the GVRP multicast MAC address. A switch that receives this frame examines the frame and extracts the VLAN IDs. GVRP then dynamically registers (creates) the VLANs and adds the receiving port to its tagged member list for the extracted VLAN IDs. The information is then transmitted out the other GVRP configured ports of the device.

Figure 3 shows an example of how VLAN Blue from end station A would be propagated across a switch network. In this figure, port 1 of Switch 4 is registered as being a member of VLAN Blue and Switch 4 declares this fact out all its ports (2 and 3) to Switch 1 and Switch 2. These two switches register this in the port egress lists of the ports (Switch 1, port 1 and Switch 2, port 1) that received the frames with the information. Switch 2, which is connected to Switch 3 and Switch 5 declares the same information to those two switches and the port egress list of each port is updated with the new information, accordingly.

Figure 3 Example of VLAN Propagation Using GVRP

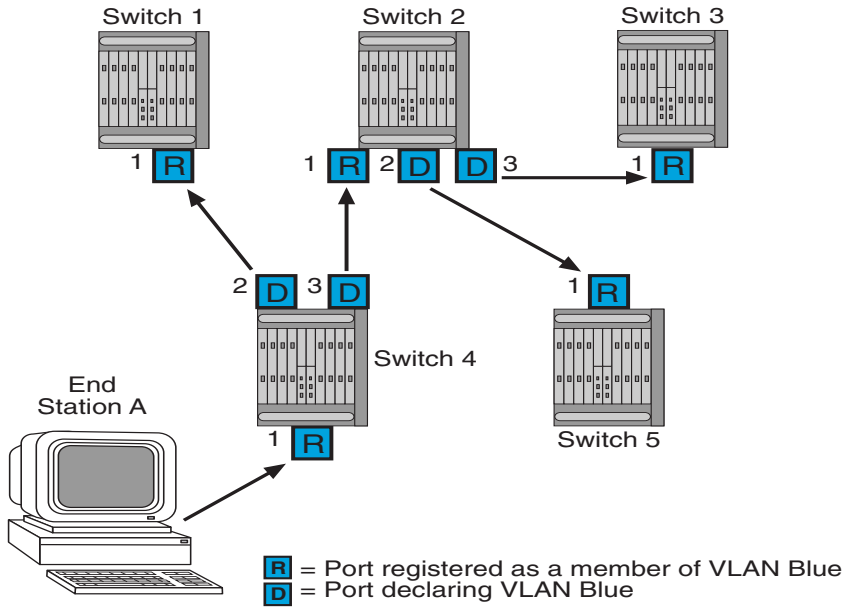

Note: If a port is set to "forbidden" for the egress list of a VLAN, then the VLAN's egress list will not be dynamically updated with that port.

Administratively configuring a VLAN on an 802.1Q switch creates a static VLAN entry that will always remain registered and will not time out. However, GVRP-created dynamic entries will time out, and their registrations will be removed from the member list if the end station is removed. This ensures that, if switches are disconnected or if end stations are removed, the registered information remains accurate.

The end result of GVRP dynamic VLAN configuration is that each port's egress list is updated with information about VLANs that reside on that port, even if the actual station on the VLAN is several hops away.

Configuring VLANs

Once you have planned your implementation strategy as described in"Preparing for VLAN Configuration" on page 3, you can begin configuring VLANs as described in this section. The following information for configuring VLANs on an Enterasys switch will be covered:

Platform Specific Differences

Matrix X Platform Configuration

The configuration of VLANs on the Matrix X-Series platform is very similar to the configuration of VLANs on the Matrix N-Series, SecureStack, and secure switch platforms, with one major exception. By default, physical ports on the Matrix X-Series are configured to route traffic not switch traffic, as is the case for the Matrix N-Series platform. Therefore, by default, no ports reside on the egress list for VLAN 1 unless the port is explicitly configured to switch traffic using the **set port mode** <***port-string***> **switched** command, and explicitly configured on VLAN 1's egress list using the **set vlan egress** <***vid***> <***port-string***> command as described in "Configuring Static VLANs" on page 11.

VLAN Naming Convention for IP Interfaces

A VLAN is identified by its ID, which is a number from 1-4094. On a Matrix X device, a VLAN entity configured on a routing interface can be specified in CLI commands in the format: **vlan.*instance*.*vlan_id***, where ***instance*** is the bridging instance, and ***vlan_id*** is the VLAN ID (1-4094). The Matrix X currently supports only one bridging instance. Therefore, ***instance*** is always 1. So, for example, to display information about VLAN 100, in either switch or router modes, you would enter:

```
show interface vlan.1.100
```

This convention is different from other Enterasys switch platforms, where the format in this instance would be **vlan** ***vlan_id***,

VLAN Constraint

VLAN constraints is a Matrix N-Series platform feature that controls the filtering database to which VLANs are allowed to belong. This feature is not supported on Matrix X, SecureStack, or secure switch platforms.

IGMP VLAN Snooping

IGMP Layer 2 snooping is a Matrix N-Series and Matrix X platform feature which allows the switch for a specific VLAN to actively participate in IGMP traffic forwarding. IGMP snooping depends on the presence of an upstream IGMP querier. Whenever it receives an IGMP query, the

switch forwards the query out the appropriate VLAN ports. IGMP snooping allows per-port traffic patterns in VLANs with multiple ports. It is disabled by default.

For more information, refer to "Configuring IGMP VLAN Snooping" on page 17.

Protected Ports

Protected Ports is a feature supported on the SecureStack and secure switch platforms that is used to prevent ports from forwarding traffic to each other, even when they are on the same VLAN. Ports can be designated as either protected or unprotected. Ports are unprotected by default. Multiple groups of protected ports are supported.

Ports that are configured to be protected:

- Cannot forward traffic to other protected ports in the same group, regardless of having the same VLAN membership.
- Can forward traffic to ports which are unprotected (not listed in any group).
- Can forward traffic to protected ports in a different group, if they are in the same VLAN.

Unprotected ports can forward traffic to both protected and unprotected ports. A port may belong to only one group of protected ports.

This feature only applies to ports within a switch. It does not apply across multiple switches in a network. Also, it is not supported on Matrix N or Matrix X platforms.

Default Settings

Table 1 lists VLAN parameters and their default values.

Table 1 Default VLAN Parameters

Parameter	Description	Default Value
garp timer	Configures the three GARP timers. The setting is critical and should only be done by someone familiar with the 802.1Q standard.	• Join timer: 20 centiseconds • Leave timer: 60 centiseconds • Leaveall timer: 1000 centiseconds
gvrp	Enables or disables the GARP VLAN Registration Protocol (GVRP) on a specific set of ports or all ports. GVRP must be enabled to allow creation of dynamic VLANs.	• Disabled at the port level • Enabled at the global level
IGMP last member query interval *(Applies to Matrix N and X only.)*	Configures the last member query interval. This is the maximum response time inserted into group-specific queries which are sent in response to Leave Group messages. It is also the amount of time between group-specific query messages.	10 seconds
IGMP VLAN max response time *(Applies to Matrix N and X only.)*	Configures the maximum query response time (in tenths of a second).	100 deciseconds

Table 1 Default VLAN Parameters (continued)

Parameter	Description	Default Value
IGMP VLAN query interval *(Applies to Matrix N and X only.)*	Configures the frequency (in seconds) of host-query frame transmissions.	125 seconds
IGMP VLAN robustness *(Applies to Matrix N and X only.)*	Configures the robustness value.	2
IGMP VLAN version *(Applies to Matrix N and X only.)*	Selects the IGMP version. Options are version 1 or version 2.	Version 2
port discard	Ports can be set to discard frames based on whether or not they contain a VLAN tag.	No frames are discarded
port ingress filter	When enabled on a port, the VLAN IDs of incoming frames are compared to the port's egress list. If the received VLAN ID does not match a VLAN ID on the port's egress list, the frame is dropped.	Enabled
port vlan ID (PVID)	802.1Q VLAN/port association.	VLAN1/ Default VLAN
protected port *(Applies to SecureStack/ Switch only)*	Prevents ports from forwarding traffic to each other, even when they are on the same VLAN.	Unprotected
vlan constraint *(Applies to Matrix N only)*	Configures VLANs to use an independent or shared filtering database.	VLANs use an independent filtering database
vlan dynamicegress	Enables or disables dynamic egress processing for a given VLAN.	Disabled
vlan egress	Configures the egress ports for a VLAN and the type of egress for the ports. Egress type can be tagged, untagged, or forbidden.	Tagged
vlan name	Associates a text name to one or more VLANs.	None

Configuring Static VLANs

Procedure 1 describes how to create and configure a static VLAN. Unspecified parameters use their default values.

Procedure 1 Static VLAN Configuration

Step	Task	Command(s)
1.	Show existing VLANs.	**`show vlan`**
2.	*(Applies to Matrix X only.)* Define the ports to be used for switched traffic.	**`set port mode`** *`port-string`* **`switched`**

Procedure 1 Static VLAN Configuration (continued)

Step	Task	Command(s)
3.	Create VLAN. Valid values are **1**–**4094**. Each *vlan-id* must be unique. If an existing *vlan-id* is entered, the existing VLAN is modified.	**set vlan create** *vlan-id*
4.	Optionally, assign a name to the VLAN. Valid strings are from 1 to 32 characters.	**set vlan name** *vlan-id string*
5.	Assign switched ports to the VLAN. This sets the port VLAN ID (PVID). The PVID determines the VLAN to which all untagged frames received on the port will be classified.	**set port vlan** *port-string vlan-id*
	Note: If the VLAN specified has not already been created, the above command will create it. It will also add the VLAN to the port's egress list as untagged, and remove the default VLAN from the port's egress list. This automatically changes the existing untagged VLAN egress permission to match the new PVID value.	
6.	Configure VLAN egress, which determines which ports a frame belonging to the VLAN may be forwarded out on.	
	Static configuration: Add the port to the VLAN egress list for the device. • The default setting, **tagged**, allows the port to transmit frames for a particular VLAN. • The **untagged** setting allows the port to transmit frames without a VLAN tag. This setting is usually used to configure a port connected to an end user device. • The **forbidden** setting prevents the port from participating in the specified VLAN and ensures that any dynamic requests for the port to join the VLAN will be ignored.	**set vlan egress** *vlan-id port-string* **forbidden** | **tagged** | **untagged**
	If necessary, remove ports from the VLAN egress list. • If specified, the **forbidden** setting will be cleared from the designated ports and the ports will be reset as allowed to egress frames, if so configured by either static or dynamic means. • If **forbidden** is not specified, tagged and untagged egress settings will be cleared from the designated ports.	**clear vlan egress** *vlan-list port-string* [**forbidden**]
	Dynamic configuration: By default, dynamic egress is disabled on all VLANs. If dynamic egress is enabled for a VLAN, the device will add the port receiving a frame to the VLAN's egress list as untagged according to the VLAN ID of the received frame.	**set vlan dynamicegress** *vlan-id* {**enable** | **disable**}

Procedure 1 Static VLAN Configuration (continued)

Step	Task	Command(s)
7.	*(Applies to Matrix N only.)* Optionally, set VLAN constraints to control the filtering database a VLAN will use for forwarding traffic. Filtering databases can be shared or independent. By default, filtering databases are independent.	**set vlan constraint** *vlan-id set-num* [**shared** | **independent**]
8.	Optionally, enable ingress filtering on a port to drop those incoming frames that do not have a VLAN ID that matches a VLAN ID on the port's egress list.	**set port ingress-filter** *port-string* **enable**
9.	Optionally, choose to discard tagged or untagged, (or both) frames on selected ports. Select **none** to allow all frames to pass through.	**set port discard** *port-string* {**tagged** | **untagged** | **none** | **both**}
10.	*(Applies to SecureStack and secure switch only.)* Optionally, configure protected ports. This prevents ports from forwarding traffic to each other, even when they are on the same VLAN. The **group-id** value identifies the assigned ports and can range from 0 to 2. You can also set a protected port group **name** of up to 32 characters in length.	**set port protected** *port-string group-id* **set port protected name** *group-id name*
11.	If the device supports routing, enter router configuration mode and configure an IP address on the VLAN interface	Matrix X: **router** **configure** **interface vlan.1.***vlan_id* **ip address** *ip-address/maxlen* **no shutdown** Matrix N/SecureStack/secure switch: **router** **enable** **configure terminal** **interface vlan** *vlan_id* **ip address** *ip-address ip-mask* **no shutdown**
	Note: Each VLAN interface must be configured for routing separately using the interface command shown above. To end configuration on one interface before configuring another, type **exit** at the command prompt. Enabling interface configuration mode is required for completing interface-specific configuration tasks.	

Example Configuration

The following shows an example Matrix N device configuration using the steps in Procedure 1. In this example, VLAN 100 is created and named VLANRED. Ports ge.1.2, 1.3 and 1.4 are assigned to VLAN 100 and added to its egress list. Ingress filtering is enabled for the VLAN ports and they are configured to discard untagged frames. VLAN 100 is then configured as a routing interface with an IP address of 120.20.20.24.

Note: Refer to Procedure 1 to determine which platform-specific commands may apply to your device when following this example configuration.

```
Matrix(su)->set vlan create 100
Matrix(su)->set vlan name 100 VLANRED
Matrix(su)->set port vlan ge.1.2-4 100
Matrix(su)->set vlan egress 100 ge.1.2-4 untagged
Matrix(su)->clear vlan egress 1 ge.1.2-4
Matrix(su)->set port ingress-filter ge.1.2-4 enable
Matrix(su)->set port discard ge.1.2-4 untagged
Matrix(su)->router
Matrix(su)->router>enable
Matrix(su)->router#configure terminal
Matrix(su)->router(config)#interface vlan 100
Matrix(su)->router(config-if(Vlan 100))#ip address 120.20.20.1/24
Matrix(su)->router(config-if(Vlan 100))#no shutdown
```

Creating a Secure Management VLAN

If you are configuring an Enterasys device for multiple VLANs, it may be desirable to configure a management-only VLAN. This allows a station connected to the management VLAN to manage the device. It also makes management secure by preventing configuration through ports assigned to other VLANs.

Procedure 2 provides an example of how to create a secure management VLAN. This example, which sets the new VLAN as VLAN 2, assumes the management station is attached to ge.1.1, and wants untagged frames.The process described in this section would be repeated on every device that is connected in the network to ensure that each device has a secure management VLAN.

Procedure 2 Secure Management VLAN Configuration

Step	Task	Command(s)
1.	*(Applies to Matrix X only.)* Configure the ports to be used as switch ports.	**`set port mode host.0.1; ge.1.1 2 switched`**
2.	Create a new VLAN.	**`set vlan create 2`**
3.	Set the PVID for the host port and the desired switch port to the VLAN created in Step 2.	**`set port vlan host.0.1; ge.1.1 2`**
4.	If not done automatically when executing the previous command, add the host port and desired switch port(s) to the new VLAN's egress list.	**`set vlan egress 2 host.0.1; ge.1.1 2 untagged`**

Procedure 2 Secure Management VLAN Configuration (continued)

Step	Task	Command(s)
5.	Set a private community name to assign to this VLAN for which you can configure access rights and policies.	**set snmp community private**

Note: By default, community name—which determines remote access for SNMP management—is set to **public** with read-write access. For more information, refer to your device's SNMP documentation.

Configuring Dynamic VLANs

Procedure 3 describes how to enable the GARP (Generic Attribute Registration Protocol) VLAN Registration Protocol (GVRP), which is needed to create dynamic VLANs. By default, GVRP is enabled globally but disabled at the port level. GVRP must be globally enabled and also enabled on specific ports in order to generate and process GVRP advertisement frames.

Note: Refer to "GARP VLAN Registration Protocol (GVRP) Support" on page 7 for conceptual information about GVRP.

Procedure 3 Dynamic VLAN Configuration

Step	Task	Command(s)
1.	Show existing GVRP configuration for a port or list of ports. If no *port-string* is entered, the global GVRP configuration and all port GVRP configurations are displayed.	**show gvrp** [***port-string***]
2.	If necessary, enable GVRP on those ports assigned to a VLAN. You must specifically enable GVRP on ports, since it is disabled on ports by default.	**set gvrp enable** *port-string*
3.	Display the existing GARP timer values.	**show garp timer** [***port-string***]
4.	Optionally, set the GARP join, leave, and leaveall timer values. Each timer value is in centiseconds.	**set garp timer** {[**join** *timer-value*] [**leave** *timer-value*] [**leaveall** *timer-value*]} *port-string*

Caution: The setting of GARP timers is critical and should only be changed by personnel familiar with 802.1Q standards.

Configuring Protocol-Based VLAN Classification

Protocol-based VLANs can be configured using the policy classification CLI commands, as shown in this section, or NetSight Policy Manager.

Procedure 4 describes how to define protocol-based frame filtering policies to assign frames to particular VLANs. Refer to your Enterasys policy configuration and CLI documentation for more information.

Note: Depending on your Enterasys switching device, your options for configuring policy classification may differ from the examples provided in this section. Refer to your device's documentation for a list of CLI commands and functions supported.

Procedure 4 Configuring Protocol-Based VLAN Classification

Step	Task	Command(s)
1.	*(Applies to Matrix X only.)* Configure the ports to be used as switch ports.	**set port mode** *port-string* **switched**
2.	Create the VLANs to which frames will be assigned by the policy. Valid values are **1**–**4094**.	**set vlan create** *vlan-id*
3.	Configure VLAN egress, which determines which ports a frame belonging to the VLAN may be forwarded out on. The default setting, **tagged**, allows the port to transmit frames for a particular VLAN.	**set vlan egress** *vlan-id port-string* [**forbidden** | **tagged** | **untagged**]
4.	Disable ingress filtering on the ingress ports on which the policy will be applied.	**set port ingress-filter** *port-string* **disable**
5.	Create the policy profile that enables PVID override. This function allows a policy rule classifying a frame to a VLAN to override PVID assignment configured with the **set port vlan** command. When none of its associated classification rules match, the configuration of the policy profile itself will determine how frames are handled by default. In this case, the default VLAN is specified with the **pvid** *pvid* parameter.	**set policy profile** *profile-index* [**name** *name*] [**pvid-status** {**enable** | **disable**}] [**pvid** *pvid*]
6.	Configure the administrative rules that will assign the policy profile to all frames received on the desired ingress ports.	**set policy rule admin-profile port** *port-string* [**port-string** *port-string*] [**admin-pid** *admin-pid*]
7.	Configure the classification rules that will define the protocol to filter on and the VLAN ID to which matching frames will be assigned.	**set policy rule** *profile-index* {*protocol data* [**mask** *mask*]} [**vlan** *vlan*]

Example Configuration

The following shows an example Matrix N device configuration using the steps in Procedure 4. This example configures a policy that ensures that IP traffic received on the specified ingress ports will be mapped to VLAN 2, while all other types of traffic will be mapped to VLAN 3.

1. Two VLANs are created: VLAN 2 and VLAN 3.
2. Ports 1 through 5 on the Gigabit Ethernet IOM in slot 4 are configured as egress ports for the VLANs while ports 8 through 10 on the Gigabit Ethernet IOM in slot 5 are configured as ingress ports that will do the policy classification.
3. Policy profile number 1 is created that enables PVID override and defines the default behavior (classify to VLAN 3) if none of the classification rules created for the profile are matched.
4. Administrative rules are created that apply policy profile number 1 to all frames received on the ingress ports ge.5.8 through 10.

5. Classification rules are created for policy profile number 1 that assign IP frames to VLAN 2. The rules identify IP frames by using the **ether** protocol parameter, which classifies on the Type field in the headers of Layer 2 Ethernet II frames, and the protocol data of 0x0800 (IP type), 0x0806 (ARP type), and 0x8035 (RARP type).

```
Matrix(su)->set vlan create 2, 3
Matrix(su)->set vlan egress 2 ge.4.1-2
Matrix(su)->set vlan egress 3 ge.4.3-5
Matrix(su)->set port ingress-filter ge.5.8-10 disable
Matrix(su)->set policy profile 1 name protocol_based_vlan pvid-status enable
   pvid 3
Matrix(su)->set policy rule admin-profile port ge.5.8 port-string ge.5.8
   admin-pid 1
Matrix(su)->set policy rule admin-profile port ge.5.9 port-string ge.5.9
   admin-pid 1
Matrix(su)->set policy rule admin-profile port ge.5.10 port-string ge.5.10
   admin-pid 1
Matrix(su)->set policy rule 1 ether 0x0800 mask 16 vlan 2
Matrix(su)->set policy rule 1 ether 0x0806 mask 16 vlan 2
Matrix(su)->set policy rule 1 ether 0x8035 mask 16 vlan 2
```

Configuring IGMP VLAN Snooping

Note: This feature is not supported on SecureStack or secure switch platforms.

IGMP Layer 2 snooping allows the Enterasys switch for a specific VLAN to actively participate in IGMP traffic forwarding. IGMP snooping depends on the presence of an upstream IGMP querier. Whenever it receives an IGMP query, the switch forwards the query out the appropriate VLAN ports. IGMP snooping allows per-port traffic patterns in VLANs with multiple ports. It is disabled by default.

For more information, refer to your device's IGMP documentation.

Procedure 5 describes how to configure IGMP snooping for a VLAN.

Procedure 5 IGMP Snooping for a VLAN Configuration

Step	Task	Command(s)
1.	Enable IGMP snooping for a VLAN or a range of VLANs.	**set igmp enable** *vlan-id*
2.	Enable querying on this VLAN, and specify the IGMP querier source address.	**set igmp query-enable** *vlan-id* **address** *ip-address*
3.	Set the version of IGMP to use. Enter **1** for IGMPV1, or **2** for IGMPV2.	**set igmp config** *vlan-id* **igmp-version 1 | 2**
4.	Set the Last Member interval value, which can be 1–255.	**set igmp config** *vlan-id* **last-member-interval** *value*
5.	Set the Max Response Time which can be 1–255 seconds.	**set igmp config** *vlan-id* **max-response-time** *seconds*
6.	Set the Query Interval, which can be 1–65535 seconds.	**set igmp config** *vlan-id* **query-interval** *seconds*

Procedure 5 IGMP Snooping for a VLAN Configuration (continued)

Step	Task	Command(s)
7.	Set the Robustness value, which can be 2–255.	**set igmp config** *vlan-id* **robustness** *value*
8.	Optionally, create a static IGMP entry, or add ports to an existing entry. The entry can be in the form of an IP multicast address or IP group address.	***set*** **igmp** *add-static {IP-multicast-address | IP-group-address vlan-id}* [**modify**] *port-string*

Monitoring VLANs

Table 2 describes the **show** commands that display information about VLAN configurations. Refer to your device's CLI documentation for a description of the output of each show command.

Table 2 Displaying VLAN Information

Task	Command
Display all existing VLANs.	**show vlan**
(Applies to Matrix N only) Display the VLAN constraint setting.	**show vlan constraint** [*vlan id*]
Display the VLAN dynamic egress setting.	**show vlan dynamicegress** [*vlan id*]
Display all static VLANs.	**show vlan static**
Display ports assigned to VLANs.	**show port vlan** [*port-string*]
Display existing GVRP settings.	**show gvrp** [*port-string*]
(*Applies to Matrix N and X only).* Display IGMP VLAN configuration.	**show igmp config** [*vlan id*]
(*Applies to Matrix N and X only).* Display IGMP enable state of VLAN.	**show igmp enable** [*vlan id*]
(*Applies to Matrix N and X only).* Display all groups on a given VLAN.	**show igmp groups** [*vlan id*]
(*Applies to Matrix N and X only).* Display IGMP VLAN query state.	**show igmp query** [*vlan id*]
Display static ports on the given vid, group.	**show igmp static** [*vlan id*]
(Applies to SecureStack/Switch only) Display port(s) configured in protected mode	**show port protected** [*port-string*] | [*group-id*]
(Applies to SecureStack/Switch only) Display the name of a specific group of protected ports.	**show port protected name** *group-id*

Terms and Definitions

Table 3 lists terms and definitions used in VLAN configuration.

Table 3 VLAN Terms and Definitions

Term	Definition
Default VLAN	The VLAN to which all ports are assigned upon initialization. The default VLAN has a VLAN ID of 1 and cannot be deleted or renamed.
Filtering Database	A database structure within the switch that keeps track of the associations between MAC addresses, VLANs, and interface (port) numbers. The Filtering Database is referred to when a switch makes a forwarding decision on a frame.
Filtering Database Identifier (FID)	Addressing information that the device learns about a VLAN is stored in the filtering database assigned to that VLAN. Several VLANs can be assigned to the same FID to allow those VLANs to share addressing information. This enables the devices in the different VLANs to communicate with each other when the individual ports have been configured to allow communication to occur. The configuration is accomplished using the Local Management VLAN Forwarding Configuration screen. By default a VLAN is assigned to the FID that matches its VLAN ID.
Forwarding List	A list of the ports on a particular device that are eligible to transmit frames for a selected VLAN.
GARP Multicast Registration Protocol (GMRP)	A GARP application that functions in a similar fashion as GVRP, except that GMRP registers multicast addresses on ports to control the flooding of multicast frames.
GARP VLAN Registration Protocol (GVRP)	A GARP application used to dynamically create VLANs across a switched network.
Generic Attribute Registration Protocol (GARP)	GARP is a protocol used to propagate state information throughout a switched network.
Port VLAN List	A per port list of all eligible VLANs whose frames can be forwarded out one specific port and the frame format (tagged or untagged) of transmissions for that port. The Port VLAN List specifies what VLANs are associated with a single port for frame transmission purposes.
Tag Header (VLAN Tag)	Four bytes of data inserted in a frame that identifies the VLAN/frame classification. The Tag Header is inserted into the frame directly after the Source MAC address field. Twelve bits of the Tag Header represent the VLAN ID. The remaining bits are other control information.
Tagged Frame	A data frame that contains a Tag Header. A VLAN aware device can add the Tag Header to any frame it transmits.
Untagged Frame	A data frame that does not have a Tag Header.
VLAN ID	A unique number (between 1 and 4094) that identifies a particular VLAN.
VLAN Name	A 32-character alphanumeric name associated with a VLAN ID. The VLAN Name is intended to make user-defined VLANs easier to identify and remember.

Revision History

Date	Description
2-1-08	New document
02-20-08	Corrected product naming conventions.

Enterasys Networks, Inc.
50 Minuteman Road
Andover, MA 01810