

Understanding Bridging and VLANs on EX Series Switches

Network switches use Layer 2 bridging protocols to discover the topology of their LAN and to forward traffic toward destinations on the LAN.

This topic explains the following concepts regarding bridging and VLANs on Juniper Networks EX Series Ethernet Switches:

- Ethernet LANs, Transparent Bridging, and VLANs on page 1
- How Bridging Works on page 2
- Types of Switch Ports on page 3
- IEEE 802.1Q Encapsulation and Tags on page 3
- Assignment of Traffic to VLANs on page 4
- Ethernet Switching Tables on page 4
- Layer 2 and Layer 3 Forwarding of VLAN Traffic on page 5
- GVRP on page 5
- Routed VLAN Interface on page 5

Ethernet LANs, Transparent Bridging, and VLANs

Ethernet is a data link layer technology, as defined by Layer 2 of the Open Systems Interconnection (OSI) model of communications protocols. Ethernet was first standardized by the IEEE in 1982, in IEEE 802.3. Ethernet is used to create LANs. The network devices, called *nodes*, on the LAN transmit data in bundles that are generally called frames or packets.

Each node on a LAN has a unique identifier so that it can be unambiguously located on the network. Ethernet uses the Layer 2 media access control (MAC) address for this purpose. MAC addresses are hardware addresses that are programmed (“burned”) into the Ethernet processor in the node.

A characteristic of Ethernet is that nodes on a LAN can transmit data frames at any time. However, the physical connecting cable between the nodes—either coaxial, copper-based (Category 5), or optical cable—can carry only a single stream of data at a time. One result of this design is that when two nodes transmit at the same time, their frames can collide on the cable and generate an error. Ethernet uses a protocol called carrier-sense multiple access with collision detection (CSMA/CD) to detect frame collisions. If a node receives a collision error message, it stops transmitting immediately and waits for a period of time before trying to send the frame again. If the node continues to detect collisions, it progressively increases the time between retransmissions in an attempt to find a time when no other data is being transmitted on the LAN. The node uses a backoff algorithm to calculate the increasing retransmission time intervals.

Ethernet LANs were originally implemented for small, simple networks that carried primarily text. Over time, LANs have become larger and more complex; the type of data they carry has grown to include voice, graphics, and video; and the increased

speed of Ethernet interfaces on LANs has resulted in exponential increases in traffic on the network.

The IEEE 802.1D-2004 standard addresses some of the problems caused by the increase in LAN and complexity. This standard defines *transparent bridging* (generally called simply bridging). Bridging divides a single physical LAN (a single *broadcast domain*) into two or more virtual LANs, or VLANs. Each *VLAN* is a collection of network nodes that are grouped together to form separate broadcast domains. On an Ethernet network that is a single LAN, all traffic is forwarded to all nodes on the LAN. On VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN. Frames that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within a VLAN and on the LAN as a whole.

On an Ethernet LAN, all network nodes must be physically connected to the same network. On VLANs, the physical location of the nodes is not important, so you can group network devices in any way that makes sense for your organization, such as by department or business function, types of network nodes, or even physical location. Each VLAN is identified by a single IP subnetwork and by standardized IEEE 802.1Q encapsulation (discussed below).

How Bridging Works

The transparent bridging protocol allows a switch to learn information about all the nodes on the LAN, including nodes on all the different VLANs. The switch uses this information to create address-lookup tables, called *Ethernet switching tables* that it consults when forwarding traffic to or toward a destination on the LAN.

Transparent bridging uses five mechanisms to create and maintain Ethernet switching tables on the switch:

- Learning
- Forwarding
- Flooding
- Filtering
- Aging

The first bridging mechanism is *learning*. When a switch is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. The switch goes through a learning process to obtain the MAC addresses of all the nodes on the network. It stores these in the Ethernet switching table. To learn MAC addresses, the switch reads all packets that it detects on the LAN or on the local VLAN, looking for MAC addresses of sending nodes. It places these addresses into its Ethernet switching table, along with two other pieces of information—the interface (or port) on which the traffic was received and the time when the address was learned.

The second bridging mechanism is *forwarding*. Switches forward traffic, passing it from an incoming interface to an outgoing interface that leads to or toward the destination. To forward frames, the switch consults the Ethernet switching table to see whether the table contains the MAC address corresponding to the frames'

destination. If the Ethernet switching table contains an entry for the desired destination address, the switch sends the traffic out the interface associated with the MAC address. The switch also consults the Ethernet switching table in the same way when transmitting frames that originate on devices connected directly to the switch. If the Ethernet switching table does not contain an entry for the desired destination address, the switch uses flooding, which is the third bridging mechanism.

Flooding is how the switch learns about destinations not in its Ethernet switching table. If this table has no entry for a particular destination MAC address, the switch floods the traffic out all interfaces except the interface on which it was received. (If traffic originates on the switch, the switch floods it out all interfaces.) When the destination node receives the flooded traffic, it sends an acknowledgment packet back to the switch, allowing it to learn the MAC address of the node and to add the address to its Ethernet switching table.

Filtering, the fourth bridging mechanism, is how broadcast traffic is limited to the local VLAN whenever possible. As the number of entries in the Ethernet switching table grows, the switch pieces together an increasingly complete picture of the VLAN and the larger LAN—of which nodes are in the local VLAN and which are on other network segments. The switch uses this information to filter traffic. Specifically, for traffic whose source and destination MAC addresses are in the local VLAN, filtering prevents the switch from forwarding this traffic to other network segments.

Finally, the switch uses *aging*, the fifth bridging mechanism, to keep the entries in the Ethernet switching table current. For each MAC address in the Ethernet switching table, the switch records a timestamp of when the information about the network node was learned. Each time the switch detects traffic from a MAC address, it updates the timestamp. A timer on the switch periodically checks the timestamp, and if it is older than a user-configured value, the switch removes the node's MAC address from the Ethernet switching table. This aging process ensures that the switch tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

Types of Switch Ports

The ports, or interfaces, on a switch operate in either access mode or trunk mode.

An interface in access mode connects to a network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The interface itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, when you boot a switch and use the factory-default configuration, or when you boot the switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode.

Trunk interfaces handle traffic for multiple VLANs, multiplexing the traffic for all those VLANs over the same physical connection. Trunk interfaces are generally used to interconnect switches to one another.

IEEE 802.1Q Encapsulation and Tags

To identify which VLAN traffic belongs to, all frames on an Ethernet VLAN are identified by a tag, as defined in the IEEE 802.1Q standard. These frames are *tagged* and are encapsulated with 802.1Q tags.

For a simple network that has only a single VLAN, all traffic has the same 802.1Q tag.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames know which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

Juniper Networks EX3200 Ethernet Switches support a maximum of 4096 VLANs. VLANs 0 and 4095 are reserved by the Juniper Networks JUNOS Software, so you cannot use them in your network.

Assignment of Traffic to VLANs

You assign traffic to a particular VLAN in one of the following ways:

- By interface (port) on the switch. You specify that all traffic received on a particular interface on the switch is assigned to a specific VLAN. If you use the default factory switch settings, all traffic received on an access interface is untagged. This traffic is part of a default VLAN, but it is not tagged with an 802.1Q tag. When configuring the switch, you specify which VLAN to assign the traffic to. You configure the VLAN either by using a VLAN number (called a VLAN ID) or by using a name, which the switch translates into a numeric VLAN ID.
- By MAC address. You can specify that all traffic received from a specific MAC address be forwarded to a specific egress interface (next hop) on the switch. This method is administratively cumbersome to configure manually, but it can be useful when you are using automated databases to manage the switches on your network.



NOTE: If a Juniper Networks EX4200 Ethernet Switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID (0 through 9) as the *slot* element of the interface name.

The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the **ethernet-switching** family. In order to include FPC 1 through FPC 9 in the default VLAN, add the **ethernet-switching** family to the configurations for those interfaces.

Ethernet Switching Tables

As EX Series switches learn the MAC addresses of the devices on local VLANs, they store them in the bridge on the switch. With each MAC address, the Ethernet switching table stores and associates the name of the interface (or port) on which the switch learned that address. The switch uses the information in this table when forwarding packets toward their destination.

Layer 2 and Layer 3 Forwarding of VLAN Traffic

To pass traffic within a VLAN, the switch uses Layer 2 forwarding protocols, including IEEE 802.1Q, Spanning Tree Protocol (STP), and GARP VLAN Registration Protocol (GVRP).

To pass traffic between two VLANs, the switch uses standard Layer 3 routing protocols, such as static routing, OSPF, and RIP. On EX Series switches, the same interfaces that support Layer 2 bridging protocols also support Layer 3 routing protocols, providing multilayer switching.

GVRP

The GARP VLAN Registration Protocol (GVRP) is an application protocol of the Generic Attribute Registration Protocol (GARP) and is defined in the IEEE 802.1Q standard. GVRP learns VLANs on a particular 802.1Q trunk port and adds the corresponding trunk port to the VLAN if the advertised VLAN is preconfigured on the switch.

The VLAN registration information sent by GVRP includes the current VLANs membership—that is, which switches are members of which VLANs—and which switch ports are in which VLAN. GVRP shares all VLAN information configured manually on a local switch.

As part of ensuring that VLAN membership information is current, GVRP removes switches and ports from the VLAN information when they become unavailable. Pruning VLAN information:

- Limits the network VLAN configuration to active participants only, reducing network overhead.
- Targets the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

Routed VLAN Interface

In a traditional network, broadcast domains consist of either physical interfaces connected to a single switch or logical interfaces connected to one or more switches through VLAN configurations. Switches send traffic to hosts that are part of the same broadcast domain, but routers are needed to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. EX Series switches use a routed VLAN interface (RVI) to perform these routing functions, using it to route data to other Layer 3 interfaces. This functionality eliminates the need for having both a switch and a router.

The RVI must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance in order for Layer 3 traffic to be routed out of it. The RVI supports IPv4, IPv6, MPLS, and IS-IS traffic. At least one Layer 2 logical interface must be operationally up in order for the RVI to be operationally up. You must configure an RVI broadcast domain or VPLS routing instance just as you would configure a VLAN on a switch. Multicast data, broadcast data, or unicast data is switched between ports within the same RVI broadcast domain or VPLS routing

instance. The RVI routes data that is destined for the switch's media access control (MAC) address.

To learn more about configuring routing protocols and policies, see the *JUNOS Software Routing Protocols Configuration Guide* at <http://www.juniper.net/techpubs/software/junos/junos95/index.html>.

- Related Topics**
- Example: Setting Up Basic Bridging and a VLAN for an EX Series Switch
 - Example: Setting Up Bridging with Multiple VLANs for EX Series Switches
 - Example: Configure Automatic VLAN Administration Using GVRP
 - Example: Connecting an Access Switch to a Distribution Switch

Published: 2009-07-22