Virtual LANs(VLAN)

VLAN Membership

When a VLAN is provided at an access layer switch, an end user must be able to gain membership to it. Two membership methods exist on Cisco Catalyst switches: static VLANs and dynamic VLANs.

. Static VLANs offer port-based membership, where switch ports are assigned to specific VLANs. End user devices become members in a VLAN based on which physical switch port they are connected to. No handshaking or unique VLAN membership protocol is needed for the end devices; they automatically assume VLAN connectivity when they connect to a port. The static port-to-VLAN membership is normally handled in hardware with application specific integrated circuits (ASICs) in the switch. This membership provides good performance because all port mappings are done at the hardware level with no complex table lookups needed.

. Dynamic VLANs are used to provide membership based on the MAC address of an end user device. When a device is connected to a switch port, the switch must query a database to establish VLAN membership. A network administrator must assign the user's MAC address to a VLAN in the database of a VLAN Membership Policy Server (VMPS). With Cisco switches, dynamic VLANs are created and managed through the use of network management tools like CiscoWorks 2000 or CiscoWorks for Switched Internetworks (CWSI). Dynamic VLANs allow a great deal of flexibility and mobility for end users, but require more administrative overhead.

Extent of VLANs

The number of VLANs that will be implemented on a network is dependent on traffic patterns, application types, segmenting common workgroups, and network management requirements. However, consideration must be given to the relationship between VLANs and the IP addressing schemes. Cisco recommends a oneto- one correspondence between VLANs and IP subnets, which means that if a Class C network address is used for a VLAN, then no more than 254 devices should be in the VLAN. Cisco also recommends that VLANs not extend beyond the Layer 2 domain of the distribution switch, i.e., the VLAN should not reach across the core of a network and into another switch block. This is designed to keep broadcasts and unnecessary movement of traffic out of the core block. VLANs can be scaled in the switch block by using two basic methods: end-to-end VLANs and local VLANs.

. End-to-end VLANs span the entire switch fabric of a network and are also called campus-wide VLANs. They are positioned to support maximum flexibility and mobility of end devices. Users are assigned to VLANs regardless of their physical location. This means that each VLAN must be made available at the access layer in every switch block. End-to-end VLANs should group users according to common requirements, following the 80/20 rule. Although only 20 percent of the traffic in a VLAN is expected to cross the network core, end-to-end VLANs make it possible for all traffic within a single VLAN to cross the core. Because all VLANs must be available at each access layer switch, VLAN trunking must be used to carry all VLANs between the access and distribution layer switches.

. In the modern network, end users require access to central resources outside their VLAN. Users must cross into the network core more frequently, making the end-to-end VLANs cumbersome and difficult to maintain. Most enterprise networks have adopted the 20/80 rule. Local VLANs deployed in this type of network. Local VLANs are designed to contain user communities based on geographic boundaries, with little regard to the amount of traffic leaving the VLAN. They range in size from a single switch in a wiring closet to an entire building. Local VLANs enables the Layer 3 function in the campus network to intelligently handle the inter-VLAN traffic loads. This provides maximum availability by using multiple paths to destinations, maximum scalability by keeping the VLAN within a switch block, and maximum manageability.

VLAN Trunking

When using VLANs in networks that have multiple interconnected switches, you need to use VLAN trunking between the switches. With VLAN trunking, the switches tag each frame sent between switches so that the receiving switch knows to what VLAN the frame belongs. End user devices connect to switch ports that provide simple connectivity to a single VLAN each. The attached devices are unaware of any VLAN structure.

A trunk link can transport more than one VLAN through a single switch port. A trunk link is not assigned to a specific VLAN.

Instead, one or more active VLANs can be transported between switches using a single physical trunk link. Connecting two switches with separate physical links for each VLAN is also possible. In addition, trunking can support multiple VLANs that have members on more than one switch. Cisco switches support two trunking protocols, namely, Inter-Switch Link (ISL) and IEEE 802.1Q.

Inter-Switch Link (ISL)

Cisco created ISL before the IEEE standardized a trunking protocol. Thus, ISL is a Cisco proprietary solution and can be used only between two Cisco switches. ISL fully encapsulates each original Ethernet frame in an ISL header and trailer. The original Ethernet frame inside the ISL header and trailer remains unchanged.

The ISL header includes a VLAN field that provides a place to encode the VLAN number. By tagging a frame with the correct VLAN number inside the header, the sending switch can ensure that the receiving switch knows to which VLAN the encapsulated frame belongs. Also, the source and destination addresses in the ISL header use MAC addresses of the sending and receiving switch, as opposed to the devices that actually sent the original frame.

802.1Q

After Cisco created ISL, the IEEE completed work on the 802.1Q standard. 802.1Q uses a different style of header to tag frames with a VLAN number than the ISL. It does not encapsulate the original frame, but adds a 4-byte header to the original Ethernet header. This additional header includes a field with which to identify the VLAN number. Because the original header has been changed, 802.1Q encapsulation forces a recalculation of the original FCS field in the Ethernet trailer, because the FCS is based on the contents of the entire frame. 802.1Q also introduces the concept of a native VLAN on a trunk. Frames belonging to this VLAN are not encapsulated with tagging information. In the event that a host is connected to an 802.1Q trunk link, that host will be able to receive and understand only the native VLAN frames.