VLAN Trunking Protocol (VTP)

Administration of network environments that consists of many interconnected switches is complicated. Cisco has developed a propriety solution to manage VLANs across such networks using the VLAN Trunking Protocol (VTP) to exchange VLAN configuration information between switches. VTP uses Layer 2 trunk frames to exchange VLAN information so that the VLAN configuration stays consistent throughout a network. VTP also manages the additions, deletions, and name changes of VLANs across multiple switches from a central point, minimizing misconfigurations and configuration inconsistencies that can cause problems, such as duplicate VLAN names or incorrect VLANtype settings.

VTP is organized into management domains or areas with common VLAN requirements. A switch can belong to only one VTP domain. Switches in different VTP domains do not share VTP information. Switches in a VTP domain advertise several attributes to their domain neighbors. Each advertisement contains information about the VTP management domain, VTP configuration revision number, known VLANs, and specific VLAN parameters.

The VTP process begins with VLAN creation on a switch called a VTP server. VTP floods advertisements throughout the VTP domain every 5 minutes, or whenever there is a change in VLAN configuration. The VTP advertisement includes a configuration revision number, VLAN names and numbers, and information about which switches have ports assigned to each VLAN. By configuring the details on one or more VTP server and propagating the information through advertisements, all switches configuration know the names and numbers of all VLANs.

The VTP Configuration Revision Number

Each time a VTP server modifies its VLAN information, it increments the configuration revision number that is sent with the VTP advertisement by 1. The VTP server then sends out a VTP advertisement that includes the new configuration revision number. When a switch receives a VTP advertisement with a larger revision number, it updates its VLAN configuration.

VTP Modes

To participate in a VTP management domain, each switch must be configured to operate in one of three modes. These modes are: server mode, client mode, and transparent mode.

Server Mode

Server mode is the default mode. In this mode, VTP servers have full control over VLAN creation and modification for their domains. All VTP information is advertised to other switches in the domain, while all received VTP information is synchronized with the other switches. Because it is the default mode, server mode can be used on any switch in a management domain, even if other server and client switches are in use. This mode provides some redundancy in the event of a server failure in the domain.

Client Mode

Client mode is a passive listening mode. Switches listens to VTP advertisements from other switches and modify their VLAN configurations accordingly. Thus the administrator is not allowed to create, change, or delete any VLANs. If other switches are in the management domain, a new switch should be configured for client mode operation. In this way, the switch will learn any existing VTP information from a server. If this switch will be used as a redundant server, it should start out in client mode to learn all VTP information from reliable sources. If the switch was initially configured for server mode instead, it might propagate incorrect information to the other domain switches. Once the switch has learned the current VTP information, it can be reconfigured for server mode.

Transparent Mode

Transparent mode does not allow the switch to participate in VTP negotiations. Thus, a switch does not advertise its own VLAN configuration, and a switch does not synchronize its VLAN database with received advertisements. VLANs can still be created, deleted, and renamed on the transparent switch. However, they will not be advertised to other neighboring switches. VTP advertisements received by a transparent switch will be forwarded on to other switches on trunk links.

VTP Pruning

A switch must forward broadcast frames out all available ports in the broadcast domain because broadcasts are destined everywhere there is a listener. Multicast frames, unless forwarded by more intelligent means, follow the same pattern. In addition, frames destined for an address that the switch has not yet learned or has forgotten must be forwarded out all ports in an attempt to find the

destination. When forwarding frames out all ports in a broadcast domain or VLAN, trunk ports are included. By default, a trunk link transports traffic from all VLANs, unless specific VLANs are removed from the trunk with the clear trunk command. In a network with several switches, trunk links are enabled between switches and VTP is used to manage the propagation of VLAN information. This causes the trunk links between switches to carry traffic from all VLANs.

VTP pruning makes more efficient use of trunk bandwidth by reducing unnecessary flooded traffic. Broadcast and unknown unicast frames on a VLAN are forwarded over a trunk link only if the switch on the receiving end of the trunk has ports in that VLAN. In other words, VTP pruning allows switches to prevent broadcasts and unknown unicasts from flowing to switches that do not have any ports in that VLAN. VTP pruning occurs as an extension to VTP version 1. When a Catalyst switch has a port associated with a VLAN, the switch sends an advertisement to its neighbor switches that it has active ports on that VLAN. The neighbors keep this information, enabling them to decide if flooded traffic from a VLAN should use a trunk port or not.

By default, VTP pruning is disabled on IOS-based and CLI-based switches. On IOS-based switches, the vtp pruning command in the VLAN database configuration mode, the can be used to enable pruning while the set vtp pruning enable command can be used to enabled VTP pruning on CLI-based switches.