Overview

This document describes how to configure Layer 2 switching features on the Cisco Industrial Ethernet 2000U Series (IE 2000U) and Connected Grid switches, hereafter referred to as *switch*.

The switch has two different types of interfaces by default: network node interfaces (NNIs) to connect to the service provider network and user network interfaces (UNIs) to connect to customer networks. Some features are supported only on one of these port types. You can also configure enhanced network interfaces (ENIs). An ENI is typically a user-network facing interface and has the same default configuration and functionality as UNIs, but can be configured to support protocol control packets for Cisco Discovery Protocol (CDP), Spanning-Tree Protocol (STP), Link Layer Discovery Protocol (LLDP), and EtherChannel Link Aggregation Control Protocol (LACP) or Port Aggregation Protocol (PAgP).

This chapter provides an overview of the following Layer 2 switching features:

VTP

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs on a network-wide basis. VTP minimizes misconfigurations and configuration inconsistencies that can cause several problems, such as duplicate VLAN names, incorrect VLAN-type specifications, and security violations.

Before you create VLANs, you must decide whether to use VTP in your network. Using VTP, you can make configuration changes centrally on one or more switches and have those changes automatically communicated to all the other switches in the network. Without VTP, you cannot send information about VLANs to other switches.

VTP is designed to work in an environment where updates are made on a single switch and are sent through VTP to other switches in the domain. It does not work well in a situation where multiple updates to the VLAN database occur simultaneously on switches in the same domain, which would result in an inconsistency in the VLAN database.

Related Topics

Chapter 2, “Configuring VLAN Trunking Protocol”

VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes as physical LANs, but you can group end stations even if they are not physically located on the same LAN segment. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router.

The VLAN feature on the switch provides the following:

- Support for up to 1005 VLANs for assigning users to VLANs associated with appropriate network resources, traffic patterns, and bandwidth.
- Support for VLAN IDs in the full 1 to 4094 range allowed by the 802.1Q standard.
- VLAN Query Protocol (VQP) for dynamic VLAN membership.
- 802.1Q trunking encapsulation on all ports for network moves, adds, and changes; management and control of broadcast and multicast traffic; and network security by establishing VLAN groups for high-security users and network resources.
- VLAN 1 minimization for reducing the risk of spanning-tree loops or storms by allowing VLAN 1 to be disabled on any individual VLAN trunk link. With this feature enabled, no user traffic is sent or received on the trunk. The switch CPU continues to send and receive control protocol frames.
- UNI-ENI isolated VLANs to isolate customer VLANs from VLANs of other customers on the same switch. Local switching does not occur among UNIs or ENIs on the switch that belong to the same UNI-ENI isolated VLAN.
- Port security on a PVLAN host to limit the number of MAC addresses learned on a port, or define which MAC addresses may be learned on a port.
- VLAN Flex Link Load Balancing to provide Layer 2 redundancy without requiring Spanning Tree Protocol (STP). A pair of interfaces configured as primary and backup links can load balance traffic based on VLAN.

Related Topics

Chapter 3, "Configuring VLANs"

VLAN Trunks

A trunk is a point-to-point link between one or more Ethernet switch interfaces and another networking device such as a router or a switch. Ethernet trunks carry the traffic of multiple VLANs over a single link, and you can extend the VLANs across an entire network. The switch supports the 802.1Q industry-standard trunking encapsulation.

Related Topics

Chapter 4, "Configuring VLAN Trunks"

Asymmetric VLAN Mapping

The Asymmetric VLAN mapping feature provides a method for restricting traffic on VLAN trunk ports. This feature lets you specify lists of VLANs that are allowed to forward traffic on the trunk port in the ingress direction, egress direction, or in both directions. This feature is supported on the CGS 2520 only.

This feature is useful in a utility substation environment where a VLAN trunk is connected between a Cisco CGS 2520 switch and an intelligent electronic device (IED). The trunk port on the Cisco CGS 2520 can be configured to allow ingress traffic for a given VLAN, such as generic object oriented substation events (GOOSE) messages from the IED, and the trunk port can be configured to allow traffic for specific VLAN IDs in the egress direction, allowing the IED to subscribe to GOOSE messages with those VLAN IDs. All other VLAN traffic on the trunk port can be blocked.

Related Topics

Chapter 5, "Configuring Asymmetric VLAN Mapping"

VMPS

Each time the client switch receives the MAC address of a new host, it sends a VLAN Query Protocol (VQP) query to the VLAN Membership Policy Server (VMPS). The query includes the newly seen MAC address and the port on which it was seen. When the VMPS receives this query, it searches its database for a MAC-address-to-VLAN mapping. The VMPS responds with a VLAN assignment for the port.

The switch cannot be a VMPS server but can act as a client to the VMPS and communicate with it through VQP.

Related Topics

Chapter 6, "Configuring VMPS"

Private VLANs

The private-VLAN feature addresses two problems that service providers face when using VLANs:

- Scalability: The switch supports up to 1005 active VLANs. If a service provider assigns one VLAN per customer, this limits the numbers of customers that the service provider can support.
- To enable IP routing, each VLAN is assigned a subnet address space or a block of addresses, which can waste the unused IP addresses and cause IP address management problems.

Using private VLANs addresses the scalability problem and provides IP address management benefits for service providers and Layer 2 security for customers.

Related Topics

Chapter 7, “Configuring Private VLANs”

IEEE 802.1Q Tunneling

802.1Q tunneling enables service providers to offer multiple point Layer 2 VPN services to customers.

Business customers of service providers often have specific requirements for VLAN IDs and the number of VLANs to be supported. The VLAN ranges required by different customers in the same service-provider network might overlap, and traffic of customers through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict customer configurations and could easily exceed the VLAN limit (4096) of the 802.1Q specification.

Using the 802.1Q tunneling (QinQ) feature, service providers can use a single VLAN to support customers who have multiple VLANs. Customer VLAN IDs (C-VLANs) are preserved, and traffic from different customers is segregated within the service-provider network, even when they appear to be in the same VLAN. Using 802.1Q tunneling expands VLAN space by using a VLAN-in-VLAN hierarchy and retagging the tagged packets.

Related Topics

Chapter 8, “Configuring IEEE 802.1Q Tunneling”

VLAN Mapping

VLAN mapping (or VLAN ID translation) on trunk ports connected to a customer network maps customer VLANs to service-provider VLANs. Packets entering the port are mapped to a service provider VLAN (S-VLAN) based on the port number and the original customer VLAN-ID (C-VLAN) of the packet. Because the VLAN ID is mapped to the S-VLAN on ingress, on the switch all forwarding operations are performed by using S-VLAN information and not C-VLAN information. Symmetrical mapping back to the C-VLAN occurs when packets exit the port.

Related Topics

Chapter 9, “Configuring VLAN Mapping”

Layer 2 Protocol Tunneling

Layer 2 protocol tunneling enables customers to control protocols such as BPDU, CDP, VTP, PAgP, LACP, and UDLD protocols to be tunneled across service-provider networks.

When protocol tunneling is enabled, edge switches on the inbound side of the service-provider network encapsulate Layer 2 protocol packets with a special MAC address and send them across the service-provider network. Core switches in the network do not process these packets but forward them as normal packets. Layer 2 protocol data units (PDUs) for CDP, STP, or VTP cross the service-provider network and are delivered to customer switches on the outbound side of the service-provider network. Identical packets are received by all customer ports on the same VLANs with these results:

- Users on each of a customer's sites can properly run STP, and every VLAN can build a correct spanning tree based on parameters from all sites and not just from the local site.
- CDP discovers and shows information about the other Cisco devices connected through the service-provider network.
- VTP provides consistent VLAN configuration throughout the customer network, propagating to all switches through the service provider that support VTP.

Related Topics

Chapter 10, "Configuring Layer 2 Protocol Tunneling"

STP

Spanning Tree Protocol (STP) is a Layer 2 link management protocol that provides path redundancy while preventing loops in the network. For a Layer 2 Ethernet network to function properly, only one active path can exist between any two stations. Multiple active paths among end stations cause loops in the network. If a loop exists in the network, end stations might receive duplicate messages. Switches might also learn end-station MAC addresses on multiple Layer 2 interfaces. These conditions result in an unstable network. Spanning-tree operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

The STP uses a spanning-tree algorithm to select one switch of a redundantly connected network as the root of the spanning tree. The algorithm calculates the best loop-free path through a switched Layer 2 network by assigning a role to each port based on the role of the port in the active topology.

STP is supported by default on NNIs, can be enabled on ENIs, and is not supported on UNIs. STP has these features:

- Up to 128 supported spanning-tree instances
- Per-VLAN spanning-tree plus (PVST+) for balancing load across VLANs
- Rapid PVST+ for balancing load across VLANs and providing rapid convergence of spanning-tree instances

Related Topics

Chapter 11, "Configuring STP"

MSTP

802.1s Multiple Spanning Tree Protocol (MSTP) enables multiple VLANs to be mapped to the same spanning-tree instance, thereby reducing the number of spanning-tree instances needed to support a large number of VLANs. The MSTP provides for multiple forwarding paths for data traffic and enables load balancing. It improves the fault tolerance of the network because a failure in one instance (forwarding

path) does not affect other instances (forwarding paths). The most common initial deployment of MSTP is in the backbone and distribution layers of a Layer 2 switched network. This deployment provides the highly available network required in a service-provider environment.

When the switch is in MST mode, the Rapid Spanning Tree Protocol (RSTP), which is based on 802.1w, is automatically enabled. The RSTP provides rapid convergence of the spanning tree through explicit handshaking that eliminates the 802.1D forwarding delay and quickly transitions root ports and designated ports to the forwarding state.

Related Topics

Chapter 12, “Configuring MSTP”

Optional STP Features

The following optional spanning-tree features are available in PVST+, rapid-PVST+, and MSTP modes on NNIs and ENIs where spanning tree has been enabled:

- Port Fast for eliminating the forwarding delay by enabling a spanning-tree port to immediately transition from the blocking state to the forwarding state
- Bridge protocol data unit (BPDU) guard for shutting down Port Fast-enabled ports that receive BPDUs
- BPDU filtering for preventing a Port Fast-enabled ports from sending or receiving BPDUs
- Root guard for preventing switches outside the network core from becoming the spanning-tree root
- Loop guard for preventing alternate or root port NNIs or ENIs from becoming designated ports because of a failure that leads to a unidirectional link

Related Topics

Chapter 13, “Configuring Optional Spanning-Tree Features”

REP

Resilient Ethernet Protocol (REP) is a Cisco proprietary protocol that provides an alternative to Spanning Tree Protocol (STP) to control network loops, to respond to link failures, and to improve convergence time. REP controls a group of ports connected in a segment, ensures that the segment does not create any bridging loops, and responds to link failures within the segment. REP provides a basis for constructing more complex networks and supports VLAN load balancing.

Related Topics

Chapter 14, “Configuring Resilient Ethernet Protocol”

UDLD

UDLD is a Layer 2 protocol that enables devices connected through fiber-optic or twisted-pair Ethernet cables to monitor the physical configuration of the cables and detect when a unidirectional link exists. All connected devices must support UDLD for the protocol to successfully identify and disable unidirectional links. When UDLD detects a unidirectional link, it disables the affected port and alerts you. Unidirectional links can cause a variety of problems, including spanning-tree topology loops.

Related Topics

Chapter 15, “Configuring UDLD”

Voice VLAN

The Voice VLAN feature enables access ports on the switch to carry IP voice traffic from a Cisco IP phone. Voice VLAN supports users connecting to both a Cisco IP phone and another data device, such as a PC, through the IP phone to a switch port. The voice traffic and data traffic can be treated differently with voice traffic having higher priority.

Related Topics

Chapter 16, “Configuring Voice VLAN”