

Inter-Switch Link and IEEE 802.1Q Frame Format

Document ID: 17056

Introduction

Prerequisites

- Requirements

- Components Used

- Conventions

- Background Theory

ISL Frame

- Field Descriptions

- Frame Size

IEEE 802.1Q Frame

- Field Descriptions

- Frame Size

QinQ

- Frame Size

- TPID

Related Information

Introduction

This document provides the basic information and a summary of the frame fields for Inter-Switch Link (ISL) and IEEE 802.1Q encapsulation.

Prerequisites

Requirements

Cisco recommends that you have knowledge of VLANs and trunking.

Components Used

This document is not restricted to specific software and hardware versions. Trunking capabilities are dependent on the hardware that is used. For more information on the system requirements to implement trunking on Cisco Catalyst series switches, refer to System Requirements to Implement Trunking.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Background Theory

Trunks are used to carry traffic that belongs to multiple VLANs between devices over the same link. A device can determine which VLAN the traffic belongs to by its VLAN identifier. The VLAN identifier is a tag that is encapsulated with the data. ISL and 802.1Q are two types of encapsulation that are used to carry data from multiple VLANs over trunk links.

ISL is a Cisco proprietary protocol for the interconnection of multiple switches and maintenance of VLAN

information as traffic goes between switches. ISL provides VLAN trunking capabilities while it maintains full wire-speed performance on Ethernet links in full-duplex or half-duplex mode. ISL operates in a point-to-point environment and can support up to 1000 VLANs. In ISL, the original frame is encapsulated and an additional header is added before the frame is carried over a trunk link. At the receiving end, the header is removed and the frame is forwarded to the assigned VLAN. ISL uses Per VLAN Spanning Tree (PVST), which runs one instance of Spanning Tree Protocol (STP) per VLAN. PVST allows the optimization of root switch placement for each VLAN and supports the load balancing of VLANs over multiple trunk links.

802.1Q is the IEEE standard for tagging frames on a trunk and supports up to 4096 VLANs. In 802.1Q, the trunking device inserts a 4-byte tag into the original frame and recomputes the frame check sequence (FCS) before the device sends the frame over the trunk link. At the receiving end, the tag is removed and the frame is forwarded to the assigned VLAN. 802.1Q does not tag frames on the native VLAN. It tags all other frames that are transmitted and received on the trunk. When you configure an 802.1Q trunk, you must make sure that you configure the same native VLAN on both sides of the trunk. IEEE 802.1Q defines a single instance of spanning tree that runs on the native VLAN for all the VLANs in the network. This is called Mono Spanning Tree (MST). This lacks the flexibility and load balancing capability of PVST that is available with ISL. However, PVST+ offers the capability to retain multiple spanning tree topologies with 802.1Q trunking.

For more information about the 802.1Q encapsulation, refer to the *Basic Characteristics of 802.1Q Trunking* section of *Trunking Between Catalyst 4500/4000, 5500/5000, and 6500/6000 Series Switches Using 802.1Q Encapsulation with Cisco CatOS System Software*.

For information on the configuration of ISL/802.1Q encapsulation on Cisco switches, refer to *VLAN Trunking Protocols Configuration Examples and TechNotes*.

ISL Frame

The ISL frame consists of three primary fields: the encapsulation frame (original frame), which is encapsulated by the ISL header, and the FCS at the end.

| | | |
|-------------------|----------------------------|------------|
| ISL Header | Encapsulation Frame | FCS |
|-------------------|----------------------------|------------|

This example shows the further expansion of the ISL header. The expansion includes the field acronyms and the number of bits for each field:

| | | | | | | | |
|-------------|----|------|------|----|-----|--------------|-----|
| No. of bits | 40 | 4 | 4 | 48 | 16 | 24 | 24 |
| Frame field | DA | TYPE | USER | SA | LEN | AAAA03(SNAP) | HSA |

| | | | | | | |
|-------------|------|------|-------|-----|---------------------------------------|-----|
| No. of bits | 15 | 1 | 16 | 16 | 8 to 196,600 bits (1 to 24,575 bytes) | 32 |
| Frame field | VLAN | BPDU | INDEX | RES | ENCAP FRAME | FCS |

Field Descriptions

This section provides detailed descriptions of the ISL frame fields.

DA Destination Address

The DA field of the ISL packet is a 40-bit destination address. This address is a multicast address and is set at "0x01-00-0C-00-00" or "0x03-00-0c-00-00". The first 40 bits of the DA field signal the receiver that the packet is in ISL format.

TYPE Frame Type

The TYPE field consists of a 4-bit code. The TYPE field indicates the type of frame that is encapsulated and can be used in the future to indicate alternative encapsulations. This table provides definitions of different TYPE codes:

| TYPE Code | Meaning |
|-----------|------------|
| 0000 | Ethernet |
| 0001 | Token Ring |
| 0010 | FDDI |
| 0011 | ATM |

USER User Defined Bits (TYPE Extension)

The USER field consists of a 4-bit code. The USER bits are used to extend the meaning of the TYPE field. The default USER field value is "0000". For Ethernet frames, the USER field bits "0" and "1" indicate the priority of the packet as it passes through the switch. Whenever traffic can be handled in a manner that allows it to be forwarded more quickly, the packets with this bit set should take advantage of the quick path. It is not required that such paths be provided.

| USER Code | Meaning |
|-----------|------------------|
| XX00 | Normal Priority |
| XX01 | Priority 1 |
| XX10 | Priority 2 |
| XX11 | Highest Priority |

SA Source Address

The SA field is the source address field of the ISL packet. The field should be set to the "802.3" MAC address of the switch port that transmits the frame. It is a 48-bit value. The receiving device may ignore the SA field of the frame.

LEN Length

The LEN field stores the actual packet size of the original packet as a 16-bit value. The LEN field represents the length of the packet in bytes, with the exclusion of the DA, TYPE, USER, SA, LEN, and FCS fields. The total length of the excluded fields is 18 bytes, so the LEN field represents the total length minus 18 bytes.

AAAA03 (SNAP) Subnetwork Access Protocol (SNAP) and Logical Link Control (LLC)

The AAAA03 SNAP field is a 24-bit constant value of "0xAAAA03".

HSA High Bits of Source Address

The HSA field is a 24-bit value. This field represents the upper 3 bytes (the manufacturer ID portion) of the SA field. The field must contain the value "0x00-00-0C".

VLAN Destination Virtual LAN ID

The VLAN field is the VLAN ID of the packet. It is a 15-bit value that is used to distinguish frames on different VLANs. This field is often referred to as the "color" of the frame.

BPDU Bridge Protocol Data Unit (BPDU) and Cisco Discovery Protocol (CDP) Indicator

The bit in the BPDU field is set for all BPDU packets that are encapsulated by the ISL frame. The BPDUs are used by the spanning tree algorithm in order to determine information about the topology of the network. This bit is also set for CDP and VLAN Trunk Protocol (VTP) frames that are encapsulated.

INDX Index

The INDX field indicates the port index of the source of the packet as it exits the switch. This field is used for diagnostic purposes only, and may be set to any value by other devices. It is a 16-bit value and is ignored in received packets.

RES Reserved for Token Ring and FDDI

The RES field is a 16-bit value. This field is used when Token Ring or FDDI packets are encapsulated with an ISL frame. In the case of Token Ring frames, the Access Control (AC) and Frame Control (FC) fields are placed here. In the case of FDDI, the FC field is placed in the Least Significant Byte (LSB) of this field. For example, an FC of "0x12" has a RES field of "0x0012". For Ethernet packets, the RES field should be set to all zeros.

ENCAP FRAME Encapsulated Frame

The ENCAP FRAME field is the encapsulated data packet, which includes its own cyclic redundancy check (CRC) value, completely unmodified. The internal frame must have a CRC value that is valid after the ISL encapsulation fields are removed. The length of this field can be from 1 to 24,575 bytes in order to accommodate Ethernet, Token Ring, and FDDI frames. A receiving switch may strip off the ISL encapsulation fields and use this ENCAP FRAME field as the frame is received (associating the appropriate VLAN and other values with the received frame as indicated for switching purposes).

FCS Frame Check Sequence

The FCS field consists of 4 bytes. This sequence contains a 32-bit CRC value, which is created by the sending MAC and is recalculated by the receiving MAC in order to check for damaged frames. The FCS is generated over the DA, SA, Length/Type, and Data fields. When an ISL header is attached, a new FCS is calculated over the entire ISL packet and added to the end of the frame.

Note: The addition of the new FCS does not alter the original FCS that is contained within the encapsulated frame.

Frame Size

The ISL frame encapsulation is 30 bytes, and the minimum FDDI packet is 17 bytes. Therefore, the minimum ISL encapsulated packet for FDDI is 47 bytes. The maximum Token Ring packet is 18,000 bytes. Therefore, the maximum ISL packet is 18,000 plus 30 bytes of ISL header, for a total of 18,030 bytes. If only Ethernet

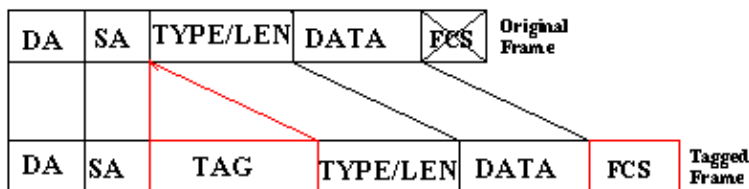
packets are encapsulated, the range of ISL frame sizes is from 94 to 1548 bytes.

The biggest implication for systems that use ISL encapsulation is that the encapsulation is a total of 30 bytes, and fragmentation is not required. Therefore, if the encapsulated packet is 1518 bytes long, the ISL packet is 1548 bytes long for Ethernet. Additionally, if packets other than Ethernet packets are encapsulated, the maximum length can be greatly increased. You must consider this length change when you evaluate whether a topology can support ISL packets size.

Another system implication is that ISL packets contain two FCSs. The first FCS is calculated for the original data. The second FCS is calculated after the packet has been encapsulated in ISL. If the original data does not contain a valid CRC, the invalid CRC is not detected until the ISL header is stripped off and the end device checks the original data FCS. This typically is not a problem for switching hardware, but can be difficult for routers and network interface cards (NICs).

IEEE 802.1Q Frame

IEEE 802.1Q uses an internal tagging mechanism which inserts a 4-byte tag field in the original Ethernet frame itself between the Source Address and Type/Length fields. Because the frame is altered, the trunking device recomputes the FCS on the modified frame.



This example shows the further expansion of the Tag field. The expansion includes the field acronyms and the number of bits for each field.

| | | | | |
|-------------|------|----------|-----|-----|
| No. of bits | 16 | 3 | 1 | 12 |
| Frame field | TPID | PRIORITY | CFI | VID |

Field Descriptions

This section provides detailed descriptions of the 802.1Q frame fields.

TPID Tag Protocol Identifier

The Tag Protocol Identifier is a 16-bit field. It is set to a value of 0x8100 in order to identify the frame as an IEEE 802.1Q-tagged frame.

Priority

Also known as user priority, this 3-bit field refers to the IEEE 802.1p priority. The field indicates the frame priority level which can be used for the prioritization of traffic. The field can represent 8 levels (0 through 7).

CFI Canonical Format Indicator

The Canonical Format Indicator is a 1-bit field. If the value of this field is 1, the MAC address is in noncanonical format. If the value is 0, the MAC address is in canonical format.

VID VLAN Identifier

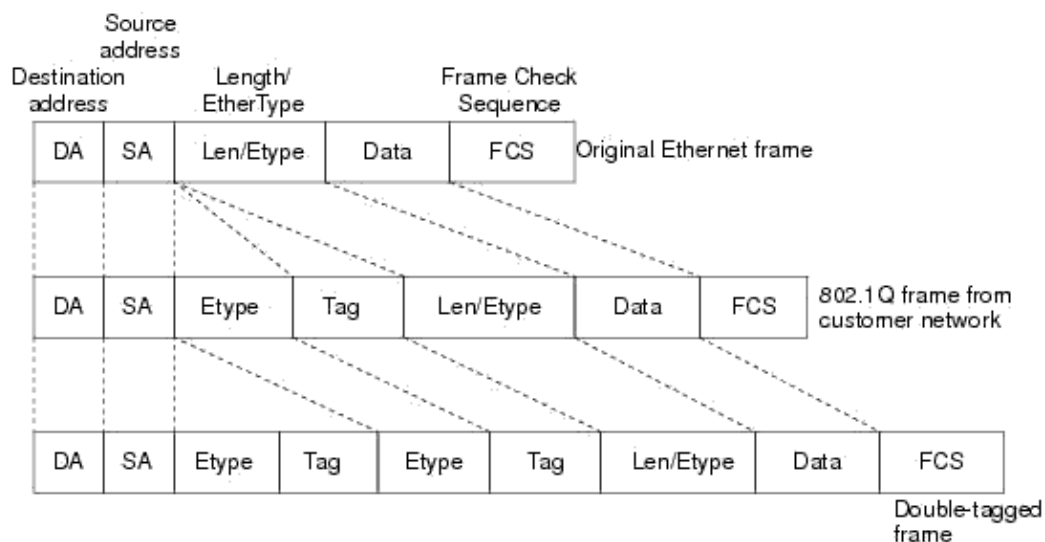
The VLAN Identifier is a 12-bit field. It uniquely identifies the VLAN to which the frame belongs. The field can have a value between 0 and 4095.

Frame Size

The 802.1Q tag is 4 bytes. Therefore, the resulting Ethernet frame can be as large as 1522 bytes. The minimum size of the Ethernet frame with 802.1Q tagging is 68 bytes.

QinQ

The QinQ Support feature adds another layer of IEEE 802.1Q tag (called "metro tag" or "PE-VLAN") to the 802.1Q tagged packets that enter the network. The purpose is to expand the VLAN space by tagging the tagged packets, thus producing a "double-tagged" frame. The expanded VLAN space allows the service provider to provide certain services, such as Internet access on specific VLANs for specific customers, yet still allows the service provider to provide other types of services for their other customers on other VLANs.



Frame Size

The default maximum transmission unit (MTU) of an interface is 1500 bytes. With an outer VLAN tag attached to an Ethernet frame, the packet size increases by 4 bytes. Therefore, it is advisable that you appropriately increase the MTU of each interface on the provider network. The recommended minimum MTU is 1504 bytes.

TPID

The QinQ frame contains the modified tag protocol identifier (TPID) value of VLAN Tags. By default, the VLAN tag uses the TPID field to identify the protocol type of the tag. The value of this field, as defined in IEEE 802.1Q, is 0x8100.

The device determines whether a received frame carries a service provider VLAN tag or a customer VLAN tag by checking the corresponding TPID value. After receiving a frame, the device compares the configured TPID value with the value of the TPID field in the frame. If the two match, the frame carries the corresponding VLAN tag. For example, if a frame carries VLAN tags with the TPID values of 0x9100 and 0x8100, respectively, while the configured TPID value of the service provider VLAN tag is 0x9100 and that of the VLAN tag for a customer network is 0x8200, the device considers that the frame carries only the service provider VLAN tag but not the customer VLAN tag.

In addition, the systems of different vendors might set the TPID of the outer VLAN tag of QinQ frames to different values. For compatibility with these systems, you can modify the TPID value so that the QinQ frames, when sent to the public network, carry the TPID value identical to the value of a particular vendor to allow interoperability with the devices of that vendor. The TPID in an Ethernet frame has the same position with the protocol type field in a frame without a VLAN tag. In order to avoid problems in packet forwarding and handling in the network, you cannot set the TPID value to any of the values in this table:

| Protocol type | Value |
|---------------|---------------|
| ARP | 0x0806 |
| PUP | 0x0200 |
| RARP | 0x8035 |
| IP | 0x0800 |
| IPv6 | 0x86DD |
| PPPoE | 0x8863/0x8864 |
| MPLS | 0x8847/0x8848 |
| IS-IS | 0x8000 |
| LACP | 0x8809 |
| 802.1x | 0x888E |

The QinQ Support feature is generally supported on whatever Cisco IOS features or protocols are supported. For example, if you can run PPPoE on the subinterface, you can configure a double-tagged frame for PPPoE. IPoQinQ supports IP packets that are double-tagged for QinQ VLAN tag termination by forwarding IP traffic with the double-tagged (also known as stacked) 802.1Q headers.

Related Information

- [System Requirements to Implement Trunking](#)
- [VLAN Trunking Protocols Configuration Examples and TechNotes](#)
- [VLAN Trunking Protocols Technology Support Page](#)
- [LAN Product Support Pages](#)
- [LAN Switching Support Page](#)
- [Technical Support & Documentation – Cisco Systems](#)