

Understanding EtherChannel Load Balancing and Redundancy on Catalyst Switches

Document ID: 12023

Contents

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Load Balancing: How to Determine the Link Across Which to Send Traffic

- Catalyst 6500/6000 Series
- Catalyst 5500/5000 Series
- Catalyst 4500/4000 Series
- Catalyst 2900XL/3500XL Series
- Catalyst 3750/3560
- Catalyst 2950/2955/3550
- Catalyst 1900/2820
- Catalyst 2948G-L3/4908G-L3 and Catalyst 8500
- Matrix of Load Balancing Methods

What Is PAgP and Where Do You Use It?

ISL/802.1Q Trunking Support on EtherChannel


Cisco Support Community - Featured Conversations

Related Information

Introduction

Fast EtherChannel allows multiple physical Fast Ethernet links to combine into one logical channel. This allows load sharing of traffic among the links in the channel as well as redundancy in the event that one or more links in the channel fail. Fast EtherChannel can be used to interconnect LAN switches, routers, servers, and clients via unshielded twisted pair (UTP) wiring or single-mode and multimode fiber. This document refers to Fast EtherChannel, Gigabit EtherChannel, port channel, channel and port group with a single term, EtherChannel. The information in the document applies to all of these EtherChannels.

This document presents the concept of load balancing and redundancy on Cisco Catalyst switches with the use of the EtherChannel. This document also covers the Port Aggregation Protocol (PAgP) and trunking support over EtherChannel. This document does not cover how to configure EtherChannel on Catalyst switches. Refer to the documents in the [Related Information](#) section for details on how to configure EtherChannel on Catalyst switches.



Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

Refer to [Cisco Technical Tips Conventions](#) for more information on document conventions.

Load Balancing: How to Determine the Link Across Which to Send Traffic

This section explains the load balance of packets over an EtherChannel group for each switch platform.

Catalyst 6500/6000 Series

EtherChannel reduces part of the binary pattern that the addresses in the frame form to a numerical value that selects one of the links in the channel in order to distribute frames across the links in a channel. EtherChannel frame distribution uses a Cisco-proprietary hashing algorithm. The algorithm is deterministic; if you use the same addresses and session information, you always hash to the same port in the channel. This method prevents out-of-order packet delivery.

Catalyst OS

In Catalyst 6500/6000 switches that run Catalyst OS (CatOS), EtherChannel aggregates the bandwidth of up to eight compatibly configured ports into a single logical link. With software releases 6.2(1) and earlier, the six- and nine-slot Catalyst 6500 series switches support a maximum of 128 EtherChannels. In software release 6.2(2) and later releases, the spanning tree feature handles the port ID. Therefore, the maximum number of EtherChannels with support is 126 for a six- or nine-slot chassis and 63 for a 13-slot chassis. All Ethernet ports on all modules, which include those on a standby Supervisor Engine, support EtherChannel with no requirement that ports be contiguous or on the same module. All ports in each EtherChannel must be the same speed. You can base the load-balance policy (frame distribution) on a MAC address (Layer 2 [L2]), an IP address (Layer 3 [L3]), or a port number (Layer 4 [L4]). You can activate these policies, respectively, if you issue the **set port channel all distribution {ip | mac} session | ip-vlan-session} [source | destination | both]** command. The **session** keyword is supported on the Supervisor Engine 2 and Supervisor Engine 720. The **ip-vlan-session** keyword is only supported on the the Supervisor Engine 720. Use this keyword in order to specify the frame distribution method, with the IP address, VLAN, and Layer 4 traffic.

If a packet does not belong to a category selection, the next-lower level category is considered. If the hardware cannot support the frame distribution method that you have selected, a `Feature not supported` error message displays.

The Cisco-proprietary hash algorithm computes a value in the range 0 to 7. With this value as a basis, a particular port in the EtherChannel is chosen. The port setup includes a mask which indicates which values the port accepts for transmission. With the maximum number of ports in a single EtherChannel, which is eight ports, each port accepts only one value. If you have four ports in the EtherChannel, each port accepts two values, and so forth. This table lists the ratios of the values that each port accepts, which depends on the number of ports in the EtherChannel:

Number of Ports in the EtherChannel	Load Balancing
8	1:1:1:1:1:1:1:1
7	2:1:1:1:1:1:1
6	2:2:1:1:1:1
5	2:2:2:1:1
4	2:2:2:2
3	3:3:2
2	4:4

Note: This table only lists the number of values, which the hash algorithm calculates, that a particular port accepts. You cannot control the port that a particular flow uses. You can only influence the load balance with a frame distribution method that results in the greatest variety.

Note: The hash algorithm cannot be configured or changed to load balance the traffic among the ports in an EtherChannel.

Note: The same Cisco-proprietary hash algorithm is also implemented in Cisco Catalyst 6500/6000 Series Switches that run Cisco IOS software.

Hence, in essence, you can only achieve perfect load balancing, even with random addresses, if you have two, four, or eight ports in the port channel.

Issue the **show port channel mod/port info** command in order to check the frame distribution policy. In version 6.1(x) and later, you can determine the port for use in the port channel to forward traffic, with the frame distribution policy as the basis. The command for this determination is **show channel hash channel-id {src_ip_addr | dest_ip_addr | src_mac_addr | dest_mac_addr | src_port | dest_port}**

[*dest_ip_addr* | *dest_mac_addr* | *dest_port*].

These are some examples:

```
1.      Console> (enable) show channel hash 865 10.10.10.1 10.10.10.2
      Selected channel port: 1/1
```

```
2.      Console> (enable) show channel hash 865 00-02-fc-26-24-94
      00-d0-c0-d7-2d-d4
```

!--- This command should be on one line.

```
Selected channel port: 1/2
```

Cisco IOS

Catalyst 6500/6000 switches that run Cisco IOS[®] system software support a maximum of 64 EtherChannels. You can form an EtherChannel with up to eight compatibly configured LAN ports on any module in a Catalyst 6500/6000 series switch. All LAN ports in each EtherChannel must be the same speed, and you must configure all the ports as either Layer 2 or Layer 3 LAN ports.

Catalyst 6500/6000 switches that run Cisco IOS system software uses a Cisco-proprietary hash algorithm. This algorithm is explained in the [Catalyst OS](#) section.

EtherChannel load balancing can use MAC addresses, IP addresses, or Layer 4 port numbers with a Policy Feature Card 2 (PFC2) and either source mode, destination mode, or both. The mode you select applies to all EtherChannels that you configure on the switch. Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel only goes to a single MAC address, use of the destination MAC address results in the choice of the same link in the channel each time. Use of source addresses or IP addresses can result in a better load balance. Issue the **port-channel load-balance {src-mac | dst-mac | src-dst-mac | src-ip | dst-ip | src-dst-ip | src-port | dst-port | src-dst-port | mpls}** global configuration command in order to configure the load balancing.

Issue the **show etherchannel load-balance** command in order to check the frame distribution policy. You can determine which interface in the EtherChannel forwards traffic, with the frame distribution policy as a basis. Issue the **remote login switch** command to log in remotely to the Switch Processor (SP) console in order to make this determination. Then, issue the **test etherchannel load-balance interface port-channel number {ip | l4port | mac} [source_ip_addr | source_mac_addr | source_l4_port] [dest_ip_addr | dest_mac_addr | dest_l4_port]** command.

These are some examples:

```
1.      6509#remote login switch
      Trying Switch ...
      Entering CONSOLE for Switch
      Type "^C^C^C" to end this session

      6509-sp#test etherchannel load-balance interface port-channel 1
      ip 10.10.10.2 10.10.10.1
```

!--- This command should be on one line.

```
Would select Gi6/1 of Po1
```

```
6509-sp#
```

```
2.      6509#remote login switch
      Trying Switch ...
      Entering CONSOLE for Switch
      Type "^C^C^C" to end this session

      6509-sp#test etherchannel load-balance interface port-channel 1 mac
      00d0.c0d7.2dd4 0002.fc26.2494
```

!--- This command should be on one line.

```
Would select Gi6/1 of Po1
```

6509-sp#

Restrictions

This section contains usage guidelines, restrictions, and troubleshooting information that apply to EtherChannel:

1. The WS-X6548-GE-TX, WS-X6548V-GE-TX, WS-X6148-GE-TX, and WS-X6148V-GE-TX modules have a limitation with EtherChannel. EtherChannel is supported on these modules for all configurations (10, 100, and 1000 Mbps speeds) but be aware of these cases of oversubscription when you configure these modules:
 - a. On these modules there is a single 1-Gigabit Ethernet uplink from the port ASIC that supports eight ports. For EtherChannel, the data from all links in a bundle goes to the port ASIC, even though the data is destined for another link. This data consumes bandwidth in the 1-Gigabit Ethernet link. For these modules, the sum total of all data on an EtherChannel cannot exceed 1 Gigabit.

You receive a message on the maximum throughput when you add a port of this module to EtherChannel.

```
C6500> (enable) set port channel 3/5,4/5 mode on
Adding a WS-X6148-GE-TX port to a channel limits the channel's
bandwidth to a maximum of 1Gig throughput
Port(s) 3/5,4/5 channel mode set to on.
C6500> (enable)
```

- b. You can also run into the oversubscription problem if you have four WS-X6148-GE-TX or WS-X6148V-GE-TX modules that run at 100 Mbps with 48 EtherChannels, and each channel has 4 ports (1 port per module).
 - c. If you use the Switch Fabric Module with the WS-X6548-GE-TX or WS-X6548V-GE-TX modules, that configuration avoids the oversubscription problem. The Switch Fabric Module interface filters and distributes the packets to the correct module per the EtherChannel bundle hash. But, you must have one port per module in the bundle. Once you have more than one port of a WS-X6548-GE-TX or WS-X6548V-GE-TX module in an EtherChannel bundle it starts to oversubscribe.

Note: With Catalyst OS software release 8.2(1), due to firmware enhancements, these oversubscription problems are no longer an issue with the WS-X6548-GE-TX and WS-X6548V-GE-TX modules.

Refer to [Ethernet and Gigabit Ethernet Switching Modules - Catalyst 6500 Series Switch Module Guide](#) for the list of other 10/100/1000 Ethernet Switching modules and Gigabit Ethernet Switching modules.

2. When you enable UplinkFast, the EtherChannel port path cost, which you set with the **set channel cost** command, for a 4-port 10/100 EtherChannel is less than the port path cost of a parallel Gigabit Ethernet link. This situation causes the slower 4-port EtherChannel link to forward and the Gigabit Ethernet link to block. Workaround is to explicitly configure a higher cost for the channel after you enable UplinkFast. Cisco bug ID [CSCds22895](#) (registered customers only) tracks this issue.
3. The WS-X6148A-GE-TX switching module has 48 oversubscribed ports in six groups of eight ports each:
 - o Ports 1, 2, 3, 4, 5, 6, 7, 8
 - o Ports 9, 10, 11, 12, 13, 14, 15, 16
 - o Ports 17, 18, 19, 20, 21, 22, 23, 24
 - o Ports 25, 26, 27, 28, 29, 30, 31, 32
 - o Ports 33, 34, 35, 36, 37, 38, 39, 40
 - o Ports 41, 42, 43, 44, 45, 46, 47, 48

The eight ports within each group use common circuitry that effectively multiplexes the group into a single, nonblocking, full-duplex Gigabit Ethernet connection to the internal switch fabric.

For each group of eight ports, the frames that are received are buffered and sent to the common Gigabit Ethernet link to the internal switch fabric. If the amount of data received for a port begins to exceed buffer capacity, flow control sends pause frames to the remote port to temporarily stop traffic and prevent frame loss.

If the frames received on any group exceeds the bandwidth of 1 Gbps, the device starts to drop the frames. These drops are not obvious as they are dropped at the internal ASIC rather than the actual interfaces. This can lead to slow throughput of packets across the device

If more throughput is required, either use ports from a line module that does not use oversubscription, or use ports from different port-groupings on oversubscribed line modules. For example, if the line module has 48 ports in groups of eight, you can select ports 1, 9, 17, 25, 33, and 41 for the same port channel.

Cisco recommends that you use 61xx, 63xx, and 64xx modules for access layer as they are usually oversubscribed and always have only a bus connection to the backplane switching bus. For a backbone connection, it is recommended that you use 65xx or 67xx modules which always have either an eight or 20 GB dedicated connection to the switch fabric.

Catalyst 5500/5000 Series

A Catalyst 5500/5000 series switch allows from two to four links to be present per Fast EtherChannel. A connection across a Fast EtherChannel is determined by source/destination address pairs. An XOR mathematical operation is performed on the last two bits of the source MAC address and the destination MAC address. This operation yields one of four results: (0 0), (0 1), (1 0), or (1 1). Each of these values points to a link in the Fast EtherChannel bundle. In the case of a two-port Fast EtherChannel, only a single bit is used in the XOR operation. This scenario yields two possible results, and each points to a link in the bundle. Circumstances can occur where one address in the source/destination pair is a constant. For example, the destination might be a server or, even more likely, a router. In that case, you still see statistical load balancing because the source address is always different. In Cisco IOS Software Release 3.1.1 and later, spanning tree is supported. As far as spanning tree is concerned, a Fast EtherChannel looks like a single bridge port, and bridge protocol data units (BPDUs) are sent down on only one of the links. A Fast EtherChannel that is in blocking mode blocks all ports on that EtherChannel connection.

EtherChannel distributes frames across the links in a channel based on the low-order bits of the source and destination MAC addresses of each frame. The frame distribution method is not configurable.

Catalyst 4500/4000 Series

Catalyst OS

In the Catalyst 4500/4000 series switches with CatOS (Supervisor Engine I and II), you can form an EtherChannel with up to eight compatibly configured Fast Ethernet or Gigabit Ethernet ports on the switch. The exact EtherChannel formation depends on the hardware. Because the spanning tree feature handles the port ID, the maximum number of channels is 126 for a six-slot chassis. In addition, you can configure an EtherChannel with the use of ports from multiple modules in CatOS release 5.x and later. All ports in an EtherChannel must be the same speed.

Catalyst OS for Catalyst 4500/4000 uses MAC address based load balancing. EtherChannel distributes frames across the links in a channel based on the low-order bits of the source and destination MAC addresses of each frame. The frame distribution method is not configurable.

Cisco IOS

A Catalyst 4500/4000 series switch with Cisco IOS Software (Supervisor Engine II+ and later) supports a maximum of 64 EtherChannels. You can form an EtherChannel with up to eight compatibly configured Ethernet interfaces on any module and across modules. All interfaces in each EtherChannel must be the same speed, and you must configure all the interfaces as either Layer 2 or Layer 3 interfaces.

EtherChannel reduces part of the binary pattern that is formed from the addresses in the frame to a numerical value that selects one of the links in the channel in order to balance the traffic load across the links in a channel. EtherChannel load balancing can use MAC addresses, IP addresses, or Layer 4 port numbers and either source mode, destination mode, or both. Use the option that provides the greatest variety in your configuration. For example, if the traffic on a channel only goes to a single MAC address, use of the destination MAC address results in the choice of the same link in the channel each time. Use of source or IP addresses can result in a better load balance. Issue the **port-channel load-balance {src-mac | dst-mac | src-dst-mac | src-ip | dst-ip | src-dst-ip | src-port | dst-port | src-dst-port}** global configuration command in order to configure load balancing. Load Balance must be configured globally and the load balancing option cannot be changed on a per port basis.

Note: The switch uses the lower order bits of source MAC address and destination MAC address in order to determine which links must be used to transmit the data. So if the data is received from the same source, then same link of the EtherChannel is used in order to forward the data.

Catalyst 2900XL/3500XL Series

A Catalyst 2900XL that runs a Cisco IOS software release that is earlier than Cisco IOS Software Release 11.2(8)SA3 chooses a link in the channel based on the link on which the destination MAC address was last heard. The software dynamically reallocates this address to another link in the channel if the link on which the address was learned is busier than the others. You can configure a Catalyst 2900XL that runs Cisco IOS Software Release 11.2(8)SA3 or later and a Catalyst 3500XL that runs Cisco IOS Software Release 11.2(8)SA6 or later in order to choose a link to be sent across the Fast EtherChannel. The switch chooses the link on the basis of the destination or source MAC address of the frame. The default is to use the source MAC address. This default means that all packets that the switch receives on a non-Fast EtherChannel port with the same MAC source address that have a destination of the MAC addresses on the other side of the channel take the same link in the channel. Use source-based forwarding when many stations that are attached to the Catalyst 2900XL/3500XL send to a few stations, such as a single router, on the other side of the Fast EtherChannel. The use of source-based forwarding in this situation evenly distributes traffic across all links in the channel. Also, the Catalyst 2900XL/3500XL switches maintain a notion of a default port on which to transmit traffic, such as Spanning Tree Protocol (STP), multicasts, and unknown unicasts.

Catalyst 3750/3560

The Catalyst 3750/3560 series switch can support up to eight compatibly configured Ethernet interfaces in an EtherChannel. The EtherChannel provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 8 Gbps (Gigabit EtherChannel) between your switch and another switch or host. With Cisco IOS Software Release 12.2(20)SE and earlier, the number of EtherChannels has a limit of 12. With Cisco IOS Software Release 12.2(25)SE and later, the number of EtherChannels has a limit of 48.

EtherChannel balances the traffic load across the links in a channel through the reduction of part of the binary pattern that the addresses in the frame form to a numerical value that selects one of the links in the channel. EtherChannel load balancing can use MAC addresses or IP addresses, source or destination addresses, or both source and destination addresses. The mode applies to all EtherChannels that are configured on the switch. You configure the load balancing and forwarding method with use of the **port-channel load-balance {dst-ip | dst-mac | src-dst-ip | src-dst-mac | src-ip | src-mac}** global configuration command.

You can find out which interface is used in the EtherChannel to forward traffic based on the load balancing method. The command for this determination is **test etherchannel load-balance interface port-channel number {ip | mac} [source_ip_add | source_mac_add] [dest_ip_add | dest_mac_add]**.

Catalyst 2950/2955/3550

The Catalyst 2950/2955 series switch can support up to eight compatibly configured Ethernet interfaces in an EtherChannel. The EtherChannel can provide full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 2 Gbps (Gigabit EtherChannel) between your switch and another switch or host. The number of EtherChannels has the limit of six with eight ports per EtherChannel.

The Catalyst 3550 series switches support both Layer 2 and Layer 3 EtherChannel, with up to eight compatibly configured Ethernet interfaces. The EtherChannel provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 8 Gbps (Gigabit EtherChannel) between your switch and another switch or host. The limit of the number of EtherChannels is the number of ports of the same type.

For the 2950/2955/3550 series switch, EtherChannel balances the traffic load across the links in a channel by randomly associating a newly learned MAC address with one of the links in the channel. EtherChannel load balancing can use either source-MAC or destination-MAC address forwarding.

With source-MAC address forwarding, when packets are forwarded to an EtherChannel, the packets are distributed across the ports in the channel based on the source-MAC address of the incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel. With destination-MAC address forwarding, when packets are forwarded to an EtherChannel, the packets are distributed across the ports in the channel based on the destination host MAC address of the incoming packet. Therefore, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel.

For the 3550 series switch, when source-MAC address forwarding is used, load distribution based on the source and destination IP address is also enabled for routed IP traffic. All routed IP traffic chooses a port based on the source and destination IP address. Packets between two IP hosts always use the same port in the channel, and traffic between any other pair of hosts can use a different port in the channel.

Issue the **port-channel load-balance {dst-mac | src-mac}** global configuration command in order to configure the load-balance and forward method.

Note: The default port is used to transmit traffic, such as Spanning Tree Protocol (STP), multicasts, and unknown unicasts. The default port can be identified from the output of the command **show etherchannel summary** by a notation of **d**.

Catalyst 1900/2820

With the enablement of PAgP, the two possible methods of link determination are preserve order and maximize load balancing between the links on the Fast EtherChannel. The [What Is PAgP and Where Do You Use It?](#) section of this document describes PAgP. The default is to maximize load balancing. PAgP is used to negotiate the configured method with the device at the other side of the channel. If preserve order is configured, the device at the other side is instructed in order to use source-based transmissions so that the Catalyst 1900/2820 always receives packets with the same source MAC address on the same link in the channel. This is the link that the Catalyst 1900/2820 always uses to send traffic to this MAC address. If maximize load balancing is configured, PAgP tells the other side that it can distribute traffic arbitrarily, and unicast traffic is transmitted by the Catalyst 1900/2820 on the link where the source address was last seen. This provides the maximum possible load-balancing configuration. When Fast EtherChannel is configured with PAgP disabled, the switch cannot negotiate with the partner about the switch learning capability. Whether the switch preserves frame ordering depends on whether the Fast EtherChannel partner performs source-based distribution. The Catalyst 1900/2820s also elect an active port. The active port is used for flooded traffic such as unknown unicast, unregistered multicast, and broadcast packets. If the port-channel mode is on (PAgP disabled), the active port is the link with the highest priority value. If the mode is desirable or auto (PAgP enabled), the active port is selected based on the priority of the links on the switch that has the higher Ethernet address. When two ports on the switch with the higher Ethernet address have the same priority, the port with the lower ifIndex is selected.

Catalyst 2948G-L3/4908G-L3 and Catalyst 8500

When one link fails, all traffic that previously used that link now uses the link next to it. For example, if Link 1 fails in a bundle, traffic that previously used Link 1 before the failure now uses Link 2.

Matrix of Load Balancing Methods

This matrix consolidates the load balancing methods that this document describes:

Platform	Address Used in XOR	Source-Based?	Destination-Based?	Source-Destination-Based?	Load Balancing Method—Configurable/Fixed?
6500/6000	Layer 2, Layer 3 addresses, Layer 4 information, or MPLS information ²	Yes	Yes	Yes	Configurable
5500/5000	Layer 2 address only	—	—	Yes	Cannot change the method
4500/4000	Layer 2, Layer 3 addresses, or Layer 4 information	Yes	Yes	Yes	Configurable
2900XL/3500XL	Layer 2 address only	Yes	Yes	—	Configurable
3750/3560	Layer 2 or Layer 3 address only	Yes	Yes	Yes	Configurable
2950/2955/3550	Layer 2 address only ¹	Yes	Yes	— ¹	Configurable
1900/2820	These platforms use a special method of load balancing. See the Catalyst 1900/2820 section for details.				
8500	Layer 3 address only	—	—	Yes	Cannot change the method

¹ For the 3550 series switch, when source-MAC address forwarding is used, load distribution based on the source and destination IP address is also enabled for routed IP traffic. All routed IP traffic chooses a port based on the source and destination IP address.

² For the 6500 series switches that run Cisco IOS, MPLS layer 2 information can also be used for load balancing MPLS packets.

What Is PAgP and Where Do You Use It?

PAgP aids in the automatic creation of EtherChannel links. PAgP packets are sent between EtherChannel-capable ports in order to negotiate the formation of a channel. Some restrictions are deliberately introduced into PAgP. The restrictions are:

- PAgP does not form a bundle on ports that are configured for dynamic VLANs. PAgP requires that all ports in the channel belong to the same VLAN or are configured as trunk ports. When a bundle already exists and a VLAN of a port is modified, all ports in the bundle are modified to match that VLAN.
- PAgP does not group ports that operate at different speeds or port duplex. If speed and duplex change when a bundle exists, PAgP changes the port speed and duplex for all ports in the bundle.
- PAgP modes are off, auto, desirable, and on. Only the combinations auto-desirable, desirable-desirable, and on-on allow the formation of a channel. The device on the other side must have PAgP set to on if a device on one side of the channel does not support PAgP, such as a router.

PAgP is currently supported on these switches:

- Catalyst 4500/4000
- Catalyst 5500/5000
- Catalyst 6500/6000
- Catalyst 2940/2950/2955/3550/3560/3750
- Catalyst 1900/2820

These switches do not support PAgP:

- Catalyst 2900XL/3500XL
- Catalyst 2948G-L3/4908G-L3
- Catalyst 8500

ISL/802.1Q Trunking Support on EtherChannel

You can configure EtherChannel connections with or without Inter-Switch Link Protocol (ISL)/IEEE 802.1Q trunking. After the formation of a channel, the configuration of any port in the channel as a trunk applies the configuration to all ports in the channel. Identically configured trunk ports can be configured as an EtherChannel. You must have all ISL or all 802.1Q; you cannot mix the two. ISL/802.1Q encapsulation, if enabled, takes place independently of the source/destination load-balancing mechanism of Fast EtherChannel. The VLAN ID has no influence on the link that a packet takes. ISL/802.1Q simply enables that trunk to belong to multiple VLANs. If trunking is not enabled, all ports that are associated with the Fast EtherChannel must belong to the same VLAN.

Cisco Support Community - Featured Conversations

[Cisco Support Community](#) is a forum for you to ask and answer questions, share suggestions, and collaborate with your peers. Below are just some of the most recent and relevant conversations happening right now.

Want to see more? Join us by clicking here
EtherChannel C2950 and C3560 mark Bayliss 11 Replies 19.05.2011 08:55
EtherChannel Hash Mechanism mhiyoshi 3 Replies 12.05.2011 05:13

Ether-channel link utilization question tschwimer 3 Replies 05.01.2008 10:05
How is Fast-Etherchannel spreading... bellocarico 1 Reply 12.12.2005 06:37
High Utilization on one port of an... alfonso.cornejo 9 Replies 09.01.2009 14:24
Building EtherChannel between C6509... sawosankung 0 Replies 22.04.2010 11:41
4 wifi bridges vs EtherChannel on... r.taccon 0 Replies 21.07.2008 06:50
Issue with multiple 3750G-24T cisconoobie 6 Replies 03.12.2007 14:19
Max length StackWise Cable g.leonard 1 Reply 04.01.2006 03:46
Start A New Discussion Subscribe

Related Information

- [Configuring EtherChannels](#)
- [Configuring Fast EtherChannel and Gigabit EtherChannel](#)
- [Configuring Fast EtherChannel and Gigabit EtherChannel](#)
- [Understanding and Configuring EtherChannel](#)
- [Creating EtherChannel Port Groups Configuring the Switch Ports](#)
- [Configuring EtherChannel](#)
- [Understanding the EtherChannel Configuring the Switch Ports](#)
- [LAN Product Support](#)
- [LAN Switching Technology Support](#)
- [Technical Support & Documentation - Cisco Systems](#)

Updated: Jul 09, 2007

Document ID: 12023

[Contacts](#) | [Feedback](#) | [Help](#) | [Site Map](#)

© 1992-2010 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)