# Configuring IPS High Bandwidth Using EtherChannel Load Balancing

This guide helps you to understand and deploy the high bandwidth features available with IPS v5.1 when used in conjunction with the EtherChannel Load Balancing (ECLB) feature of Cisco Catalyst OS 8.5(1) and Cisco Catalyst IOS.

This document has three parts:

- Introduction — EtherChannel and how ECLB works in conjunction with IPS appliances.
- ECLB/IPS Deployment — step-by-step configuration and analysis of ECLB and IPS deployment.
- Catalyst IOS ECLB Deployment — step-by-step configuration and analysis of Catalyst IOS ECLB deployment.

# Introduction

## Cisco EtherChannel

Cisco EtherChannel® technology builds upon standards-based 802.3 full-duplex Fast Ethernet to provide network managers with a reliable, high-speed solution for the campus network backbone. EtherChannel technology provides bandwidth scalability within the campus by providing up to 800 Mbps, 8 Gbps, or 80 Gbps of aggregate bandwidth for a Fast EtherChannel, Gigabit EtherChannel, or 10

Gigabit EtherChannel connection, respectively. Each of these connection speeds can vary in amounts equal to the speed of the links used (100 Mbps, 1 Gbps, or 10 Gbps). Even in the most bandwidth-demanding situations, EtherChannel technology helps aggregate traffic and keep oversubscription to a minimum, while providing effective link-resiliency mechanisms. For more information on Cisco EtherChannel® technology see:

- Cisco EtherChannel Technology
- Understanding EtherChannel Load Balancing and Redundancy on Catalyst Switches Document ID: 12023

# IPS v5.1 High Bandwidth Overview

ECLB is a feature of Catalyst switch code that allows a Catalyst switch to split flows over different physical paths. With this technology you can place multiple IPS appliances in on-a-stick mode, clustered on an EtherChannel, to split load among the sensors. This deployment results in, incremental IPS bandwidth increase for each IPS appliance you add to the ether-channel. So if you start with a single IPS appliance with a throughput of 500 megabytes (MB) and then add a second sensor with the same throughput capabilities to the EtherChannel, you will have 1 gigabyte of resulting throughput.

You should understand the following about deploying ECLB:

- Although the ECLB name insinuates load balancing, it is not. It is actually load splitting. ECLB works by hashing the source and destination IP addresses of a flow to determine what physical port a flow should take when traversing the Catalyst.
- ECLB for IPS Service Modules works only with IPS v5.1 (or later) and Cisco Catalyst OS 8.5(1) (or later) on the CAT6k platform and Cisco Catalyst IOS.
- ECLB for IPS appliances works with both CatOS and Cat IOS.
- The IPS appliances must be in on-a-stick mode, meaning that the IPS appliance can only use one sensing port on that Catalyst switch. That port is trunked so that the IPS appliance has an inbound and outbound path to and from the switch.
- Up to eight ports can be defined in an EtherChannel. This means that you can add up to eight IPS appliances on a single Catalyst switch.

- Although native ECLB can use multiple hash algorithms, IPS has only been tested using the source and destination IP address as the EtherChannel Hash algorithm.
- ECLB guarantees that return traffic takes the same path as the initial traffic flow, thereby ensuring that IPS appliances have the symmetric traffic path required to work effectively.
- The IPS solution does not maintain state if a sensor goes down and that TCP flow is forced through a different IPS appliance. The resulting TCP flow must be reset before traffic can continue to flow.

# ECLB Packet Profile

This section has four parts:

ECLB IPS Basic Design Example — Explains the basic VLAN design for ECLB.

ECLB Flow #1 — Provides an initial ECLB packet flow example.

ECLB Flow #2 — Provides a subsequent ECLB packet flow example.

Value-Add of ECLB and IPS On-a-Stick - Explains the value added by ECLB and IPS on-a-stick when joined.

## ECLB IPS Basic Design Example

Figure 1 details how ECLB works.

There are two IPS appliances in this deployment. One is connected to port 2/22 on the Catalyst switch, the other is connected to port 2/23. Both are using 802.1q trunking VLANs 10 and 11. Both ports have been set to trunk the appropriate VLANs and added to the same EtherChannel.

*Figure 1 VLAN Design for ECLB*

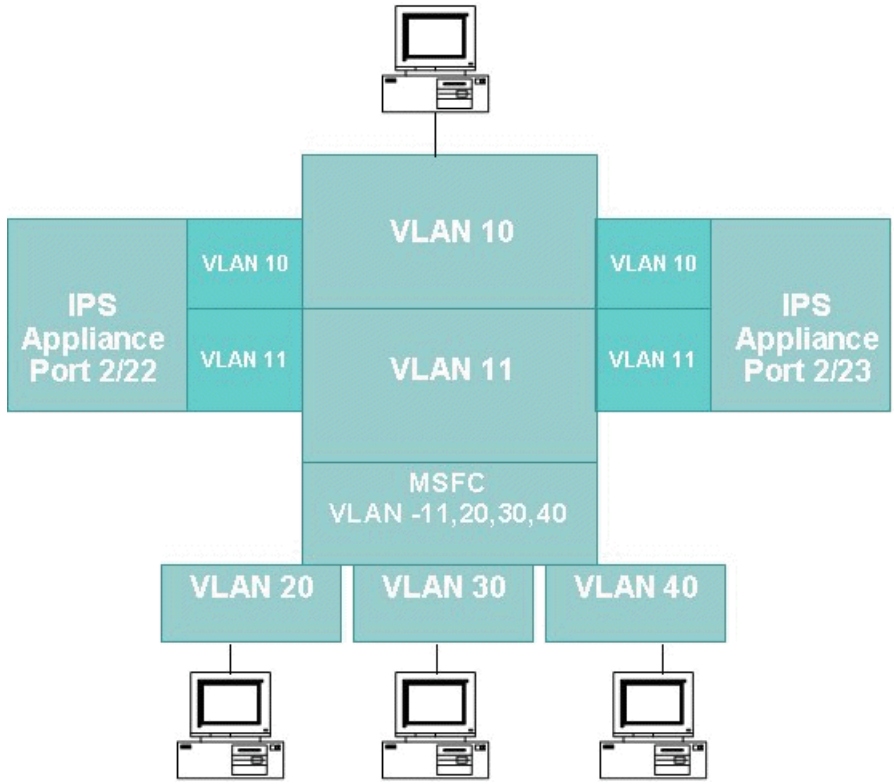


The source IP address is 10.1.10.1, which is part of VLAN 10, and the destination address is 10.1.20.1, which is part of VLAN 20.

VLAN11 is a key part of the IPS on-a-stick implementation. Since we are routing between subnets (10.1.10.x <-> 10.1.20.x) we need to define a gateway for those subnets on the MSFC. However, if you define an address on the MSFC for VLAN10, the client uses address resolution protocol (ARP) for their default gateway and the switch passes their packets directly to each other around the IPS. The traffic would never be seen or evaluated by an IPS appliance.

To get around this problem, IPS simply bridges the same subnet between two VLANS, 10 and 11. So you must define your MSFC gateway on VLAN 11. This forces traffic to go through the IPS appliance for evaluation before routing takes place.

VLAN 11 must be associated with a hardware port and that port *must* have a link. We achieve that in this example by pulling a cable into another switch port allocated just for that purpose.

For the procedure to configure this deployment, see ECLB/IPS Deployment, page 9.

## ECLB Flow #1

For the first flow, shown in Figure 2, we source traffic from VLAN 10 destined to VLAN 20. A step-by-step synopsis of how ECLB traffic passes though the switch follows:

1. The first packet with a source address in 10.1.10.1 (VLAN 10) and a destination of 10.1.20.1 (VLAN20) is generated and placed on the Catalyst back plane.
2. The ECLB engine evaluates the packet and creates a hash using the source and destination IP addresses.
3. Based on the result of the hash and the number of ports defined in ECLB, the engine decides to which port to forward the packet. In this example the path it takes is via port 2/22.
4. The IPS appliance that is connected to port 2/22 gets the packet that came from VLAN 10, so it has a VLAN 10 tag on the front of the packet.
5. The IPS on port 2/22 evaluates the traffic against it's IPS engine. If the traffic is OK, a new tag of VLAN11 is added to the packet and it is put back out the same port. If there is an alert fired, IPS appliance takes whatever action is defined for that alert, including possibly a Deny Packet action.
6. After the packet arrives in VLAN 11, the switch knows (because of ARP and ARP replies) what port the end device resides on in VLAN 11 (the MSFC's interface in VLAN 11) and forwards the packet there.
7. Traffic that has been put back out on the 2/22 port is then evaluated by the MSFC for routing.
8. The MSFC/Supervisor sees the destination subnet of 10.1.20.x and works in conjunction with the ARP tables to direct the packet to the appropriate port.

*Figure 2 ECLB Flow #1*

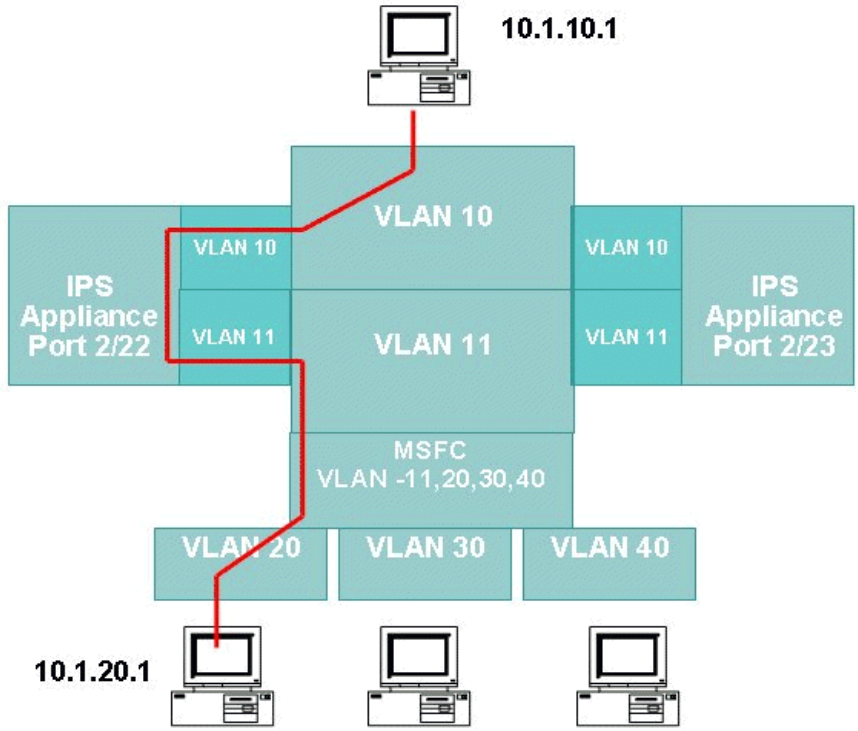


**ECLB maps a flow to a physical port within an EtherChannel; in this case port 2/22. This mapping will persist until the flow is torn down.**

## ECLB Flow #2

The traffic path for the second flow, shown in Figure 3, is almost exactly the same as the first except that, based on the ECLB calculation that another path is selected, this time the selected path includes the port to which we have connected the other IPS appliance. Therefore, we effectively split our load and increase our bandwidth using the integrated technology of ECLB and IPS on-a-stick mode.

In the second flow, we are going to set up traffic from source VLAN 10 to destination VLAN 40. A step-by-step synopsis of how ECLB traffic passes though the switch follows:

1. The first packet, with a source address in VLAN 10 10.1.10.1 and a destination of VLAN40 10.1.40.1, is generated and placed on the back plane.

- The ECLB engine evaluates the packet and creates a hash using the source and destination IP addresses.
- Based on the result of the hash and the number of ports defined in ECLB, the engine decides to which port to forward the packet. In this example, the path it takes is via 2/23, which is the second IPS appliance in the EtherChannel.
- The IPS appliance that is connected to port 2/23 gets the packet that came from VLAN 10, so it has a VLAN 10 tag on the front of the packet.
- The IPS appliance on 2/23 evaluates the traffic against it's IPS engine. If the traffic is OK, a new tag of VLAN11 is added to the packet and it is put back out the same port. If there is a alert fired, the IPS appliance takes whatever action is defined for that alert, including a drop packet.
- Traffic that has been put back out on the 2/23 port is then evaluated by the MSFC for routing.
- The MSFC sees the destination subnet of 10.1.40.x and works in conjunction with the ARP tables to route the packet to the appropriate port.

*Figure 3 ECLB Flow # 2*

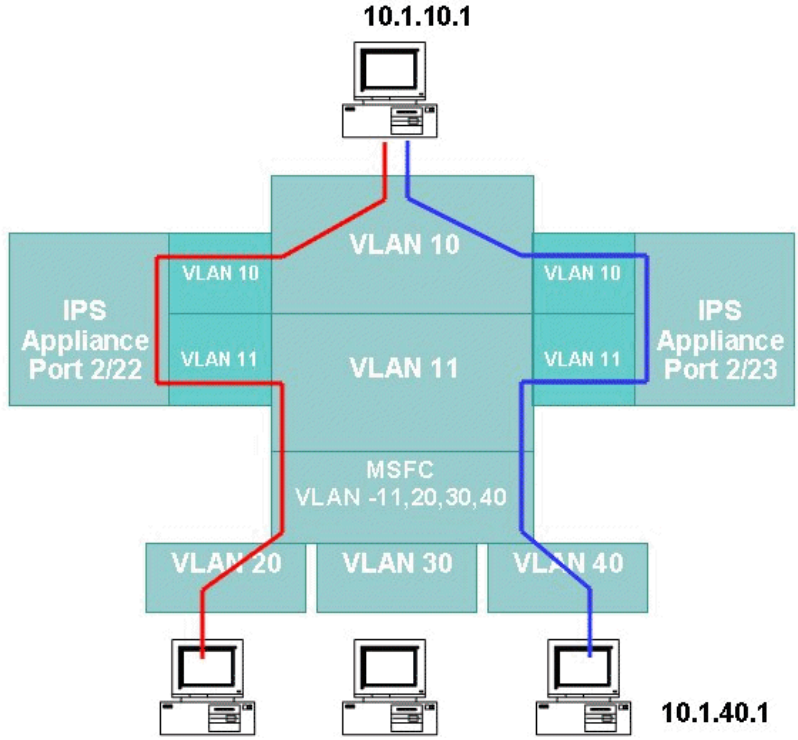


**A new connection is requested. ECLB maps a flow to a physical port within an EtherChannel; in this case port 2/23.**

You can see from the example that the load has been distributed to different paths using the bandwidth of both IPS appliances.

# Value-Add of ECLB and IPS On-a-Stick

Simply stated, incremental IPS throughput is the value-add for ECLB in conjunction with IPS on-a-stick. With the addition of ECLB you can take an existing high-end system with approximately 750 megabytes of protection and turn that into (750meg * 8) or approximately 6 Gigs of theoretical IPS throughput.

# ECLB/IPS Deployment

To properly deploy ECLB you must perform a number of basic tasks. These tasks are detailed in the following sections:

- Configure the IPS Appliances
- Configure the CAT6K - L2
- Configure the CAT6K MSFC - L3
- Test the ECLB Configurations

This section is a step-by-step guide on how to deploy the example scenario that was described in the ECLB Packet Profile section. We are going to put two IPS appliances on an EtherChannel that will bridge traffic between VLAN 10 and VLAN 11. You must configure the IPS appliance first. Otherwise layer two will not be up and the EtherChannel and trunks on the Catalyst will not come up correctly.

## Configure the IPS Appliances

**Step 1** Run SETUP on both IPS appliances.

**a.** Set the password if required.

**b.** Configure the sensor name.

**c.** Configure the sensor management address and gateway.

**d.** Configure your management access lists.

**Tip** You do not need to define any interfaces at this point.

**e.** Save the configuration

**Step 2** Launch IDS for both IPS appliances.

**a.** Enable the interface that you plan to use on each device.

**Tip** In our example, the on-a-stick interface from one IPS appliance is connected to 2/22 the other to 2/23 on the CAT6K device.

b. Configure the VLAN pairs. Use the correct interface, a sub-interface of 1, and subnets of 10 and 11.

c. Associate the VLAN pairs to the virtual sensor interface.

# Configure the CAT6K - L2

Having configured the IPS appliances, you now configure the CAT6K device.

**Step 1** Configure the following information:

a. Set port 2/31 to 10.

b. Set port 2/32 to 20.

c. Set port 2/30 to 30.

d. Set port 2/33 to 40.

e. Set port 2/31 to 10.

f. Set trunk 2/22 to none.

g. Set trunk 2/23 to none.

h. Set trunk 2/22 to noneg dot1q 10,11.

i. Set trunk 2/23 to noneg dot1q 10,11.

j. Set port channel 2/22 mode to on.

k. Set port channel 2/23 mode to on.

**Step 2** Verify that the EtherChannel has come up correctly.

a. Enter the following command:

```
Console> (enable) sh channel stat
```

b. Confirm the channel statistics report complies with Table 1.

*Table 1 Channel Statistics*

| | Port 2/22 | Port 2/23 |
|---|---|---|
| **Channel** | 1708 | 1708 |
| **Transmitted** | 0 | 0 |
| **Received** | 0 | 0 |
| **InFlush** | 0 | 0 |
| **RetnFlush** | 0 | 0 |
| **OutFlush** | 0 | 0 |
| **InError** | 0 | 0 |

# Configure the CAT6K MSFC - L3

Next, configure the IP address for interfaces VLAN 11, VLAN 20, VLAN 30, and VLAN 40.

**Step 1** Configure VLAN 11 to use 10.1.10.254 /24.

**Step 2** Configure VLAN 20 to use 10.1.20.254 /24.

**Step 3** Configure VLAN 30 to use 10.1.30.254 /24.

**Step 4** Configure VLAN 40 to use 10.1.40.254 /24.

# Test the ECLB Configurations

The next step in recreating the example is to test the ECLB configurations.

**Step 1** Do a continuous ping from the attack server to 10.1.20.1.

**Step 2** Do a continuous ping from the attack server to 10.1.30.1.

**Step 3** Do a continuous ping from the attack server to 10.1.40.1.

**Tip** You must do all three pings or the results from this lab may be very misleading.

**Step 4** Verify on the CAT6K that some traffic is traversing both port 2/22 and port 2/23.

**a.** Enter the following command:

```
Console> (enable) sh channel traff
```

**b.** Confirm the traffic reported is similar to that shown in Table 2.

***Table 2 Cat 6K Channel Traffic***

| | Port 2/22 | Port 2/23 |
|---|---|---|
| **Channel ID** | 1708 | 1708 |
| **Port** | 2/22 | 2/23 |
| **Rx-Ucst** | 63.57% | 36.43% |
| **Tx-Ucst** | 63.57% | 36.43% |
| **Rx-Mcst** | 99.66% | 0.34% |
| **Tx-Mcst** | 91.49 | 8.51% |
| **Rx-Bcst** | 0.00% | 100.00% |
| **Tx-Bcst** | 0.00% | 100.00% |

You should see approximately 33% of the traffic on one interface and 66% on the other interface.

# Catalyst IOS ECLB Deployment

This section details how to configure and test CAT IOS ECLB.

## Configure Catalyst IOS ECLB

**Step 1** Set port 2/31 to vlan 10 with the following commands:

```
interface FastEthernet2/31
no ip address
switchport
switchport access vlan 10
```

**Step 2** Set port 2/32 to vlan 20 with the following commands:

```
interface FastEthernet2/32
no ip address
switchport
switchport access vlan 20
```

**Step 3** Set port 2/30 to vlan 30 with the following commands:

```
interface FastEthernet2/30
no ip address
switchport
switchport access vlan 30
```

**Step 4** Set port to vlan 2/33 40 with the following commands:

```
interface FastEthernet2/33
 no ip address
 switchport
 switchport access vlan 40
```

**Step 5** Set trunk 2/22 noneg dot1q 10,11 and apply EtherChannel to that trunk with the following commands:

```
interface FastEthernet2/22
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,11
 switchport mode trunk
 channel-group 2 mode on
```

**Step 6** Set trunk 2/23 noneg dot1q 10,11 with the following commands:

```
interface FastEthernet2/23
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,11
 switchport mode trunk
 channel-group 2 mode on
```

**Step 7** Configure the ELCB engine to use the source and destination IP addresses for is hash calculation

```
port-channel load-balance src-dst-ip
```

# Test the CAT IOS ECLB Configurations

After you configure the CAT IOS ECLB, next you must test it.

**Step 1** Do a continuous ping from the attack server to 10.1.20.1.

**Step 2** Do a continuous ping from the attack server to 10.1.30.1.

**Step 3** Do a continuous ping from the attack server to 10.1.40.1.

**Tip** You must do all three pings or the viewing of the results from this lab may be very misleading.

**Step 4** Verify on the Catalyst IOS that traffic is traversing both ports in the EtherChannel: port Fastethernet2/22 and port Fastethernet 2/23. Enter the following commands to see which port will be used for each address pair:

```
test etherchannel load-balance interface port-channel 2 ip 10.1.10.1
10.1.20.1
test etherchannel load-balance interface port-channel 2 ip 10.1.10.1
10.1.30.1
test etherchannel load-balance interface port-channel 2 ip 10.1.10.1
10.1.40.1
```