



Configuring Flex Links and the MAC Address-Table Move Update Feature

- [Finding Feature Information, on page 1](#)
- [Restrictions for Configuring Flex Links and MAC Address-Table Move Update, on page 1](#)
- [Information About Flex Links and MAC Address-Table Move Update, on page 2](#)
- [How to Configure Flex Links and the MAC Address-Table Move Update Feature, on page 5](#)
- [Monitoring Flex Links, Multicast Fast Convergence, and MAC Address-Table Move Update, on page 9](#)
- [Configuration Examples for Flex Links, on page 10](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table at the end of this module.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

Restrictions for Configuring Flex Links and MAC Address-Table Move Update

- Flex Links are supported only on Layer 2 ports and port channels.
- You can configure up to 16 backup links.
- You can configure only one Flex Links backup link for any active link, and it must be a different interface from the active interface.
- An interface can belong to only one Flex Links pair. An interface can be a backup link for only one active link. An active link cannot belong to another Flex Links pair.

- Neither of the links can be a port that belongs to an EtherChannel. However, you can configure two port channels (EtherChannel logical interfaces) as Flex Links, and you can configure a port channel and a physical interface as Flex Links, with either the port channel or the physical interface as the active link.
- A backup link does not have to be the same type (Gigabit Ethernet or port channel) as the active link. However, you should configure both Flex Links with similar characteristics so that there are no loops or changes in behavior if the standby link begins to forward traffic.
- STP is disabled on Flex Links ports. A Flex Links port does not participate in STP, even if the VLANs present on the port are configured for STP. When STP is not enabled, be sure that there are no loops in the configured topology.

Information About Flex Links and MAC Address-Table Move Update

Flex Links

Flex Links are a pair of a Layer 2 interfaces (device ports or port channels) where one interface is configured to act as a backup to the other. The feature provides an alternative solution to the Spanning Tree Protocol (STP). Users can disable STP and still retain basic link redundancy. Flex Links are typically configured in service provider or enterprise networks where customers do not want to run STP on the device. If the device is running STP, Flex Links are not necessary because STP already provides link-level redundancy or backup.

You configure Flex Links on one Layer 2 interface (the active link) by assigning another Layer 2 interface as the Flex Links or backup link. On devices, the Flex Links can be on the same device or on another device in the stack. When one of the links is up and forwarding traffic, the other link is in standby mode, ready to begin forwarding traffic if the other link shuts down. At any given time, only one of the interfaces is in the linkup state and forwarding traffic. If the primary link shuts down, the standby link starts forwarding traffic. When the active link comes back up, it goes into standby mode and does not forward traffic. STP is disabled on Flex Links interfaces.

Flex Links Configuration

In the following figure, ports 1 and 2 on device A are connected to uplink switches B and C. Because they are configured as Flex Links, only one of the interfaces is forwarding traffic; the other is in standby mode. If port 1 is the active link, it begins forwarding traffic between port 1 and switch B; the link between port 2 (the backup link) and switch C is not forwarding traffic. If port 1 goes down, port 2 comes up and starts forwarding traffic to switch C. When port 1 comes back up, it goes into standby mode and does not forward traffic; port 2 continues forwarding traffic.

You can also configure a preemption function, specifying the preferred port for forwarding traffic. For example, you can configure the Flex Links pair with preemption mode. In the scenario shown, when port 1 comes back up and has more bandwidth than port 2, port 1 begins forwarding traffic after 60 seconds. Port 2 becomes the standby port. You do this by entering the **switchport backup interface preemption mode bandwidth** and **switchport backup interface preemption delay** interface configuration commands.

If a primary (forwarding) link goes down, a trap notifies the network management stations. If the standby link goes down, a trap notifies the users.

Flex Links are supported only on Layer 2 ports and port channels, not on VLANs or on Layer 3 ports.

VLAN Flex Links Load Balancing and Support

VLAN Flex Links load balancing allows users to configure a Flex Links pair so that both ports simultaneously forward the traffic for some mutually exclusive VLANs. For example, if Flex Links ports are configured for 1 to 100 VLANs, the traffic of the first 50 VLANs can be forwarded on one port and the rest on the other port. If one of the ports fail, the other active port forwards all the traffic. When the failed port comes back up, it resumes forwarding traffic in the preferred VLANs. In addition to providing the redundancy, this Flex Links pair can be used for load balancing. Flex Links VLAN load balancing does not impose any restrictions on uplink devices.

Multicast Fast Convergence with Flex Links Failover

Multicast fast convergence reduces the multicast traffic convergence time after a Flex Links failure. Multicast fast convergence is implemented by a combination of learning the backup link as an mrouter port, generating IGMP reports, and leaking IGMP reports.

Learning the Other Flex Links Port as the mrouter Port

In a typical multicast network, there is a querier for each VLAN. A device deployed at the edge of a network has one of its Flex Links ports receiving queries. Flex Links ports are also always forwarding at any given time.

A port that receives queries is added as an mrouter port on the device. An mrouter port is part of all the multicast groups learned by the device. After a changeover, queries are received by the other Flex Links port. The other Flex Links port is then learned as the mrouter port. After changeover, multicast traffic then flows through the other Flex Links port. To achieve faster convergence of traffic, both Flex Links ports are learned as mrouter ports whenever either Flex Links port is learned as the mrouter port. Both Flex Links ports are always part of multicast groups.

Although both Flex Links ports are part of the groups in normal operation mode, all traffic on the backup port is blocked. The normal multicast data flow is not affected by the addition of the backup port as an mrouter port. When the changeover happens, the backup port is unblocked, allowing the traffic to flow. In this case, the upstream multicast data flows as soon as the backup port is unblocked.

Generating IGMP Reports

When the backup link comes up after the changeover, the upstream new distribution device does not start forwarding multicast data, because the port on the upstream router, which is connected to the blocked Flex Links port, is not part of any multicast group. The reports for the multicast groups were not forwarded by the downstream device because the backup link is blocked. The data does not flow on this port, until it learns the multicast groups, which occurs only after it receives reports.

The reports are sent by hosts when a general query is received, and a general query is sent within 60 seconds in normal scenarios. When the backup link starts forwarding, to achieve faster convergence of multicast data, the downstream device immediately sends proxy reports for all the learned groups on this port without waiting for a general query.

Leaking IGMP Reports

To achieve multicast traffic convergence with minimal loss, a redundant data path must be set up before the Flex Links active link goes down. This can be achieved by leaking only IGMP report packets on the Flex Links backup link. These leaked IGMP report messages are processed by upstream distribution routers, so

multicast data traffic gets forwarded to the backup interface. Because all incoming traffic on the backup interface is dropped at the ingress of the access device, no duplicate multicast traffic is received by the host. When the Flex Links active link fails, the access device starts accepting traffic from the backup link immediately. The only disadvantage of this scheme is that it consumes bandwidth on the link between the distribution devices and on the backup link between the distribution and access devices. This feature is disabled by default and can be configured by using the **switchport backup interface *interface-id* multicast fast-convergence** command.

When this feature has been enabled at changeover, the device does not generate the proxy reports on the backup port, which became the forwarding port.

MAC Address-Table Move Update

The MAC address-table move update feature allows the device to provide rapid bidirectional convergence when a primary (forwarding) link goes down and the standby link begins forwarding traffic.

Flex Links VLAN Load Balancing Configuration Guidelines

- For Flex Links VLAN load balancing, you must choose the preferred VLANs on the backup interface.
- You cannot configure a preemption mechanism and VLAN load balancing for the same Flex Links pair.

MAC Address-Table Move Update Configuration Guidelines

- You can enable and configure this feature on the access device to *send* the MAC address-table move updates.
- You can enable and configure this feature on the uplink devices to *get* the MAC address-table move updates.

Default Flex Links and MAC Address-Table Move Update Configuration

- Flex Links is not configured, and there are no backup interfaces defined.
- The preemption mode is off.
- The preemption delay is 35 seconds.
- The MAC address-table move update feature is not configured on the device.

How to Configure Flex Links and the MAC Address-Table Move Update Feature

Configuring Flex Links

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: Switch# <code>configure terminal</code>	Enters global configuration mode.
Step 2	interface <i>interface-id</i> Example: Switch(conf)# <code>interface gigabitethernet1/0/1</code>	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface).
Step 3	switchport backup interface <i>interface-id</i> Example: Switch(conf-if)# <code>switchport backup interface gigabitethernet1/0/2</code>	Configures a physical Layer 2 interface (or port channel) as part of a Flex Links pair with the interface. When one link is forwarding traffic, the other interface is in standby mode.
Step 4	end Example: Switch(conf-if)# <code>end</code>	Returns to privileged EXEC mode.

Configuring a Preemption Scheme for a Pair of Flex Links

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: Switch# <code>configure terminal</code>	Enters global configuration mode

	Command or Action	Purpose
Step 2	interface <i>interface-id</i> Example: <pre>Switch(conf)# interface gigabitethernet1/0/1</pre>	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface).
Step 3	switchport backup interface <i>interface-id</i> Example: <pre>Switch(conf-if)# switchport backup interface gigabitethernet1/0/2</pre>	Configures a physical Layer 2 interface (or port channel) as part of a Flex Links pair with the interface. When one link is forwarding traffic, the other interface is in standby mode.
Step 4	switchport backup interface <i>interface-id</i> preemption mode [forced bandwidth off] Example: <pre>Switch(conf-if)# switchport backup interface gigabitethernet1/0/2 preemption mode forced</pre>	Configures a preemption mechanism and delay for a Flex Links interface pair. You can configure the preemption as: <ul style="list-style-type: none"> • forced—(Optional) The active interface always preempts the backup. • bandwidth—(Optional) The interface with the higher bandwidth always acts as the active interface. • off—(Optional) No preemption occurs from active to backup.
Step 5	switchport backup interface <i>interface-id</i> preemption delay <i>delay-time</i> Example: <pre>Switch(conf-if)# switchport backup interface gigabitethernet1/0/2 preemption delay 50</pre>	Configures the time delay until a port preempts another port. <p>Note Setting a delay time only works with forced and bandwidth modes.</p>
Step 6	end Example: <pre>Switch(conf-if)# end</pre>	Returns to privileged EXEC mode.
Step 7	show interface [<i>interface-id</i>] switchport backup Example: <pre>Switch# show interface gigabitethernet1/0/2 switchport backup</pre>	Verifies the configuration.

	Command or Action	Purpose
Step 8	copy running-config startup config Example: Switch# <code>copy running-config startup config</code>	(Optional) Saves your entries in the device startup configuration file.

Configuring VLAN Load Balancing on Flex Links

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: Switch# <code>configure terminal</code>	Enters global configuration mode.
Step 2	interface <i>interface-id</i> Example: Switch (config)# <code>interface gigabitethernet2/0/6</code>	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface).
Step 3	switchport backup interface <i>interface-id</i> prefer vlan <i>vlan-range</i> Example: Switch (config-if)# <code>switchport backup interface gigabitethernet2/0/8 prefer vlan 2</code>	Configures a physical Layer 2 interface (or port channel) as part of a Flex Links pair with the interface and specifies the VLANs carried on the interface. The VLAN ID range is 1 to 4094.
Step 4	end Example: Switch (config-if)# <code>end</code>	Returns to privileged EXEC mode.

Configuring MAC Address-Table Move Update

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: Switch# configure terminal	Enters global configuration mode.
Step 2	interface <i>interface-id</i> Example: Switch# interface gigabitethernet1/0/1	Specifies the interface, and enters interface configuration mode. The interface can be a physical Layer 2 interface or a port channel (logical interface).
Step 3	Use one of the following: <ul style="list-style-type: none"> • switchport backup interface <i>interface-id</i> • switchport backup interface <i>interface-id</i> mmu primary vlan <i>vlan-id</i> Example: Switch(config-if)# switchport backup interface gigabitethernet0/2 mmu primary vlan 2	Configures a physical Layer 2 interface (or port channel), as part of a Flex Links pair with the interface. The MAC address-table move update VLAN is the lowest VLAN ID on the interface. Configure a physical Layer 2 interface (or port channel) and specifies the VLAN ID on the interface, which is used for sending the MAC address-table move update. When one link is forwarding traffic, the other interface is in standby mode.
Step 4	end Example: Switch(config-if)# end	Returns to global configuration mode.
Step 5	mac address-table move update transmit Example: Switch(config)# mac address-table move update transmit	Enables the access device to send MAC address-table move updates to other devices in the network if the primary link goes down and the device starts forwarding traffic through the standby link. Enter command mac address-table move update on the device, for MMU packets to update MAC tables. When the primary link comes back up, the MAC tables need to reconverge and this command will transmit the MMU, that will establish the behavior.
Step 6	end Example:	Returns to privileged EXEC mode.

	Command or Action	Purpose
	Switch(config)# end	

Configuring a Device to Obtain and Process MAC Address-Table Move Update Messages

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: Switch# configure terminal	Enters global configuration mode
Step 2	mac address-table move update receive Example: Switch (config)# mac address-table move update receive	Enables the device to obtain and processes the MAC address-table move updates.
Step 3	end Example: Switch (config)# end	Returns to privileged EXEC mode.

Monitoring Flex Links, Multicast Fast Convergence, and MAC Address-Table Move Update

Command	Purpose
show interface [<i>interface-id</i>] switchport backup	Displays the Flex Links backup interface configured for an interface or all the configured Flex Links and the state of each active and backup interface (up or standby mode).
show ip igmp profile address-table move update <i>profile-id</i>	Displays the specified IGMP profile or all the IGMP profiles defined on the device.
show mac address-table move update	Displays the MAC address-table move update information on the device.

Configuration Examples for Flex Links

Configuring Flex Links: Examples

This example shows how to verify the configuration after you configure an interface with a backup interface:

```
Switch# show interface switchport backup

Switch Backup Interface Pairs:
Active Interface Backup Interface State
-----
GigabitEthernet1/0/1 GigabitEthernet1/0/2 Active Up/Backup Standby
```

This example shows how to verify the configuration after you configure the preemption mode as forced for a backup interface pair:

```
Switch# show interface switchport backup detail

Switch Backup Interface Pairs:

Active Interface Backup Interface State
-----
GigabitEthernet1/0/211 GigabitEthernet1/0/2 Active Up/Backup Standby
Interface Pair : Gi1/0/1, Gi1/0/2
Preemption Mode : forced
Preemption Delay : 50 seconds
Bandwidth : 100000 Kbit (Gi1/0/1), 100000 Kbit (Gi1/0/2)
Mac Address Move Update Vlan : auto
```

Configuring VLAN Load Balancing on Flex Links: Examples

In the following example, VLANs 1 to 50, 60, and 100 to 120 are configured on the device:

```
Switch(config)# interface gigabitethernet 2/0/6
Switch(config-if)# switchport backup interface gigabitethernet 2/0/8 prefer vlan 60,100-120
```

When both interfaces are up, Gi2/0/8 forwards traffic for VLANs 60 and 100 to 120 and Gi2/0/6 forwards traffic for VLANs 1 to 50.

```
Switch# show interfaces switchport backup

Switch Backup Interface Pairs:

Active Interface      Backup Interface      State
-----
GigabitEthernet2/0/6  GigabitEthernet2/0/8  Active Up/Backup Standby

Vlans Preferred on Active Interface: 1-50
Vlans Preferred on Backup Interface: 60, 100-120
```

When a Flex Links interface goes down (LINK_DOWN), VLANs preferred on this interface are moved to the peer interface of the Flex Links pair. In this example, if interface Gi2/0/6 goes down, Gi2/0/8 carries all VLANs of the Flex Links pair.

```
Switch# show interfaces switchport backup

Switch Backup Interface Pairs:

Active Interface      Backup Interface      State
-----
GigabitEthernet2/0/6  GigabitEthernet2/0/8  Active Down/Backup Up

Vlans Preferred on Active Interface: 1-50
Vlans Preferred on Backup Interface: 60, 100-120
```

When a Flex Links interface comes up, VLANs preferred on this interface are blocked on the peer interface and moved to the forwarding state on the interface that has just come up. In this example, if interface Gi2/0/6 comes up, VLANs preferred on this interface are blocked on the peer interface Gi2/0/8 and forwarded on Gi2/0/6.

```
Switch# show interfaces switchport backup

Switch Backup Interface Pairs:

Active Interface      Backup Interface      State
-----
GigabitEthernet2/0/6  GigabitEthernet2/0/8  Active Up/Backup Standby

Vlans Preferred on Active Interface: 1-50
Vlans Preferred on Backup Interface: 60, 100-120

Switch# show interfaces switchport backup detail

Switch Backup Interface Pairs:

Active Interface      Backup Interface      State
-----
FastEthernet1/0/3     FastEthernet1/0/4    Active Down/Backup Up

Vlans Preferred on Active Interface: 1-2,5-4094
Vlans Preferred on Backup Interface: 3-4
Preemption Mode      : off
Bandwidth : 10000 Kbit (Fa1/0/3), 100000 Kbit (Fa1/0/4)
Mac Address Move Update Vlan : auto
```

Configuring the MAC Address-Table Move Update: Examples

This example shows how to verify the configuration after you configure an access device to send MAC address-table move updates:

```
Switch# show mac address-table move update

Switch-ID : 010b.4630.1780
Dst mac-address : 0180.c200.0010
Vlans/Macs supported : 1023/8320
Default/Current settings: Rcv Off/On, Xmt Off/On
Max packets per min : Rcv 40, Xmt 60
```

```

Rcv packet count : 5
Rcv conforming packet count : 5
Rcv invalid packet count : 0
Rcv packet count this min : 0
Rcv threshold exceed count : 0
Rcv last sequence# this min : 0
Rcv last interface : Po2
Rcv last src-mac-address : 000b.462d.c502
Rcv last switch-ID : 0403.fd6a.8700
Xmt packet count : 0
Xmt packet count this min : 0
Xmt threshold exceed count : 0
Xmt pak buf unavail cnt : 0
Xmt last interface : None

```

Configuring Multicast Fast Convergence with Flex Links Failover: Examples

These are configuration examples for learning the other Flex Links port as the mrouter port when Flex Links is configured on GigabitEthernet1/0/11 and GigabitEthernet1/0/12, and output for the **show interfaces switchport backup** command:

```

Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface GigabitEthernet1/0/11
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport backup interface GigabitEthernet1/0/12
Switch(config-if)# exit
Switch(config)# interface GigabitEthernet1/0/12
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport mode trunk
Switch(config-if)# end
Switch# show interfaces switchport backup detail
Switch Backup Interface Pairs:
Active Interface Backup Interface State
GigabitEthernet1/0/11 GigabitEthernet1/0/12 Active Up/Backup Standby
Preemption Mode : off
Multicast Fast Convergence : Off
Bandwidth : 100000 Kbit (Gi1/0/11), 100000 Kbit (Gi1/0/12)
Mac Address Move Update Vlan : auto

```

This output shows a querier for VLANs 1 and 401, with their queries reaching the device through GigabitEthernet1/0/11:

```

Switch# show ip igmp snooping querier

```

Vlan	IP Address	IGMP Version	Port
1	1.1.1.1	v2	Gi1/0/11
401	41.41.41.1	v2	Gi1/0/11

This example is output for the **show ip igmp snooping mrouter** command for VLANs 1 and 401:

```

Switch# show ip igmp snooping mrouter

```

Vlan	ports
----	-----

```

1      Gi1/0/11 (dynamic), Gi1/0/12 (dynamic)
401   Gi1/0/11 (dynamic), Gi1/0/12 (dynamic)

```

Similarly, both Flex Links ports are part of learned groups. In this example, GigabitEthernet2/0/11 is a receiver/host in VLAN 1, which is interested in two multicast groups:

```
Switch# show ip igmp snooping groups
```

Vlan	Group	Type	Version	Port List
1	228.1.5.1	igmp	v2	Gi1/0/11, Gi1/0/12, Gi2/0/11
1	228.1.5.2	igmp	v2	Gi1/0/11, Gi1/0/12, Gi2/0/11

When a host responds to the general query, the device forwards this report on all the mrouter ports. In this example, when a host sends a report for the group 228.1.5.1, it is forwarded only on GigabitEthernet1/0/11, because the backup port GigabitEthernet1/0/12 is blocked. When the active link, GigabitEthernet1/0/11, goes down, the backup port, GigabitEthernet1/0/12, begins forwarding.

As soon as this port starts forwarding, the device sends proxy reports for the groups 228.1.5.1 and 228.1.5.2 on behalf of the host. The upstream router learns the groups and starts forwarding multicast data. This is the default behavior of Flex Links. This behavior changes when the user configures fast convergence using the **switchport backup interface gigabitEthernet 1/0/12 multicast fast-convergence** command. This example shows turning on this feature:

```

Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# interface gigabitEthernet 1/0/11
Switch(config-if)# switchport backup interface gigabitEthernet 1/0/12 multicast
fast-convergence
Switch(config-if)# exit
Switch# show interfaces switchport backup detail

Switch Backup Interface Pairs:
Active      Interface      Backup Interface State
-----
GigabitEthernet1/0/11 GigabitEthernet1/0/12 Active Up/Backup Standby
Preemption Mode : off
Multicast Fast Convergence : On
Bandwidth : 100000 Kbit (Gi1/0/11), 100000 Kbit (Gi1/0/12)
Mac Address Move Update Vlan : auto

```

This output shows a querier for VLAN 1 and 401 with their queries reaching the device through GigabitEthernet1/0/11:

```
Switch# show ip igmp snooping querier
```

Vlan	IP Address	IGMP Version	Port
1	1.1.1.1	v2	Gi1/0/11
401	41.41.41.1	v2	Gi1/0/11

This is output for the **show ip igmp snooping mrouter** command for VLAN 1 and 401:

```
Switch# show ip igmp snooping mrouter
```

Vlan	ports
------	-------

```

-----
1      Gi1/0/11(dynamic), Gi1/0/12(dynamic)
401    Gi1/0/11(dynamic), Gi1/0/12(dynamic)

```

Similarly, both the Flex Links ports are a part of the learned groups. In this example, GigabitEthernet2/0/11 is a receiver/host in VLAN 1, which is interested in two multicast groups:

```
Switch# show ip igmp snooping groups
```

```

Vlan  Group      Type  Version  Port List
-----
1     228.1.5.1    igmp  v2       Gi1/0/11, Gi1/0/12, Gi2/0/11
1     228.1.5.2    igmp  v2       Gi1/0/11, Gi1/0/12, Gi2/0/11

```

Whenever a host responds to the general query, the device forwards this report on all the mrouter ports. When you turn on this feature through the command-line port, and when a report is forwarded by the device on GigabitEthernet1/0/11, it is also leaked to the backup port GigabitEthernet1/0/12. The upstream router learns the groups and starts forwarding multicast data, which is dropped at the ingress because GigabitEthernet1/0/12 is blocked. When the active link, GigabitEthernet1/0/11, goes down, the backup port, GigabitEthernet1/0/12, begins forwarding. You do not need to send any proxy reports as the multicast data is already being forwarded by the upstream router. By leaking reports to the backup port, a redundant multicast path has been set up, and the time taken for the multicast traffic convergence is very minimal.