Chapter 5: VLAN Trunks and EtherChannel Bundles

Instructor Materials

CCNP Enterprise: Core Networking

Chapter 5 Content

This chapter covers the following content:

- **VLAN Trunking Protocol (VTP) -** This section provides an overview of how switches become aware of other switches and prevent forwarding loops.

- **Dynamic Trunking Protocol (DTP) -** This section examines the improvements made to STP for faster convergence.

- **EtherChannel Bundle -**This section explains how multiple physical interfaces can be combined to form a logical interface to increase throughput and provide seamless resiliency.

VLAN Trunking Protocol

- Cisco created the proprietary protocol, VLAN Trunking Protocol (VTP), to reduce the burden of provisioning VLANs on switches.
- Switches that participate in the same VTP domain can have a VLAN created once on a VTP server and propagated to other VTP client switches in the same VTP domain.

The Roles of VTP

There are four roles in the VTP architecture:

VTP Roll	Description
Server	The server switch is responsible for the creation, modification, and deletion of VLANs within the VTP domain.
Client	The client switch receives VTP advertisements and modifies the VLANs on that switch. VLANs cannot be configured locally on a VTP client.
Transparent	VTP transparent switches receive and forward VTP advertisements but do not modify the local VLAN database. VLANs are configured only locally.
Off	A switch does not participate in VTP advertisements and does not forward them out of any ports either. VLANs are configured only locally.

The Versions of VTP

There are three versions of VTP:

- Version 1 is default.
- Versions 1 and 2 have limited propagation to VLANs numbered 1 to 1005.
- VTP Version 3 allows for the full range of VLANs 1 to 4094.

VTP supports having multiple VTP servers in a domain. These servers process updates from other VTP servers just as a client does.

If a VTP domain is Version 3, the primary VTP server must be set with the executive command **vtp primary.**

VTP Communication

VTP advertises updates by using a multicast address across the trunk links for advertising updates to all the switches in the VTP domain. The three main types of VTP advertisements:

Communication Types	Description
Summary	This advertisement occurs every 300 seconds or when a VLAN is added, removed, or changed. It includes the VTP version, domain, configuration revision number, and time stamp.
Subset	This advertisement occurs after a VLAN configuration change occurs. It contains all the relevant information for the switches to make changes to the VLANs on them.
Client Requests	This advertisement is a request by a client to receive the more detailed subset advertisement. This occurs when a switch with a lower revision number joins the VTP domain and observes a summary advertisement with a higher revision than it has stored locally.

VTP Configuration

The following are the steps for configuring VTP:

Terms	Description
Step 1	Define the VTP version with the command **vtp version** {**1** | **2** | **3**}.
Step 2	Define the VTP domain with the command **vtp domain** *domain-name*. Changing the VTP domain resets the local switch's version to 0.
Step 3	Define the VTP switch role with the command **vtp mode** { **server** | **client** | **transparent** | **none** }
Step 4	(Optional) Secure the VTP domain with the command **vtp password** *password* (This step is optional but recommended because it helps prevent unauthorized switches from joining the VTP domain.)

VTP Configuration Example

Example 5-1 demonstrates the VTP configuration on SW1, SW2, SW3, and SW6.

The figure shows sample configurations for three of the VTP roles: SW1 as a client, SW3 as transparent, and the other switches as VTP clients.

Example 5-1 *Configuring the VTP Domain*

```
SW1(config)# vtp domain CiscoPress
Changing VTP domain name from CCNP to CiscoPress
SW1(config)# vtp version 3
09:08:11.965: %SW_VLAN-6-OLD_CONFIG_FILE_READ: Old version 2 VLAN configuration
  file detected and read OK. Version 3 files will be written in the future.
09:08:12.085: %SW_VLAN-6-VTP_DOMAIN_NAME_CHG: VTP domain name changed to CISCO.
SW1(config)# vtp mode server
Setting device to VTP Server mode for VLANS.
SW1(config)# vtp password PASSWORD
Setting device VTP password to PASSWORD
SW1(config)# exit
SW1# vtp primary
This system is becoming primary server for feature vlan
No conflicting VTP3 devices found.
Do you want to continue? [confirm]
09:25:02.038: %SW_VLAN-4-VTP_PRIMARY_SERVER_CHG: 0062.ec9d.c500 has become the
   primary  server for the VLAN VTP feature
```

```
SW2(config)# vtp version 3
SW2(config)# vtp domain CISCO
SW2(config)# vtp mode client
SW2(config)# vtp password PASSWORD
Setting device VTP password to PASSWORD
```

```
SW3(config)# vtp version 3
SW3(config)# vtp domain CISCO
SW3(config)# vtp mode transparent
SW3(config)# vtp password PASSWORD
```

```
SW6(config)# vtp version 3
SW6(config)# vtp domain CISCO
SW6(config)# vtp mode client
SW6(config)# vtp password PASSWORD
```

VTP Verification

The VTP status is verified with the command **show vtp status** as shown in the example.

The most important information displayed is the VTP version, VTP domain name, VTP mode, the number of VLANs (standard and extended), and the configuration version.

Example 5-2 *Verifying VTP*

```
SW1# show vtp status
VTP Version capable             : 1 to 3
VTP version running             : 3
VTP Domain Name                 : CISCO
VTP Pruning Mode                : Disabled
VTP Traps Generation            : Disabled
Device ID                       : 0062.ec9d.c500

Feature VLAN:
--------------
VTP Operating Mode              : Server
Number of existing VLANs        : 5
Number of existing extended VLANs : 0
Maximum VLANs supported locally   : 4096
Configuration Revision          : 1
Primary ID                      : 0062.ec9d.c500
Primary Description             : SW1
MD5 digest                      : 0x9D 0xE3 0xCD 0x04 0x22 0x70 0xED 0x73
                                  0x96 0xDE 0x0B 0x7A 0x15 0x65 0xE2 0x65
! The following information is used for other functions not covered in the Enterprise
! Core exam and are not directly relevant and will not be explained
Feature MST:
--------------
VTP Operating Mode              : Transparent

Feature UNKNOWN:
--------------
VTP Operating Mode              : Transparent
```

```
SW2# show vtp status | i version run|Operating|VLANS|Revision
VTP version running             : 3
VTP Operating Mode              : Client
Configuration Revision          : 1
VTP Operating Mode              : Transparent
VTP Operating Mode              : Transparent
```

VTP Verification (Cont.)

It is very important that every switch that connects to a VTP domain has the VTP revision number reset to 0. Failing to reset the revision number on a switch could result in the switch providing an update to the VTP server.

This is not an issue if VLANs are added but is catastrophic if VLANs are removed because those VLANs will be removed throughout the domain.

Example 5-3 *Creating VLANs on the VTP Domain Server*

```
SW1(config)# vlan 10
SW1(config-vlan)# name PCs
SW1(config-vlan)# vlan 20
SW1(config-vlan)# name VoIP
SW1(config-vlan)# vlan 30
SW1(config-vlan)# name Guest
```

```
SW1# show vtp status | i version run|Operating|VLANS|Revision
VTP version running             : 3
VTP Operating Mode                : Primary Server
Configuration Revision            : 4
VTP Operating Mode                : Transparent
VTP Operating Mode                : Transparent
```

```
SW6# show vlan

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Gi1/0/1, Gi1/0/2, Gi1/0/4
                                                Gi1/0/5, Gi1/0/6, Gi1/0/7
                                                Gi1/0/8, Gi1/0/9, Gi1/0/10
                                                Gi1/0/11, Gi1/0/12, Gi1/0/13
                                                Gi1/0/14, Gi1/0/15, Gi1/0/16
                                                Gi1/0/17, Gi1/0/18, Gi1/0/19
                                                Gi1/0/20, Gi1/0/21, Gi1/0/22
                                                Gi1/0/23, Gi1/0/24
10   PCs                              active
20   VoIP                             active
30   Guest                            active
1002 fddi-default                     act/unsup
1003 trcrf-default                    act/unsup
1004 fddinet-default                  act/unsup
1005 trbrf-default                    act/unsup
```

Dynamic Trunking Protocol

- Dynamic trunk ports are established by the switch port sending Dynamic Trunking Protocol (DTP) packets to negotiate whether the other end can be a trunk port.
- DTP advertises itself every 30 seconds to neighbors so that they are kept aware of its status. DTP requires that the VTP domain match between the two switches.

DTP Modes

- **Trunk -** The command **switchport mode trunk** statically places the switch port as a trunk. This mode advertises DTP packets to the other end to establish a dynamic trunk.
- **Dynamic desirable -** Using the command **switchport mode dynamic desirable,** the switch port acts as an access port, but listens for and advertises DTP packets to the other end to establish a dynamic trunk. If it is successful in negotiation, the port becomes a trunk port.
- **Dynamic auto:** Using the command **switchport mode dynamic auto,** the switch port acts as an access port, but it listens for DTP packets. It responds to DTP packets and upon successful negotiation the port becomes a trunk port.

Terms	Trunk	Dynamic Desirable	Dynamic Auto
Trunk	✓	✓	✓
Dynamic Desirable	✓	✓	✓
Dynamic Auto	✓	✓	X

DTP Mode Configuration

Example 5-5 shows the configuration of DTP on SW1's Gi1/0/2 as a dynamic auto switch port and SW2's Gi1/0/1 as a dynamic desirable switch port.

The trunk port status is verified with the command **show interface** [*interface-id*] **trunk**, as shown in Example 5-6.

Example 5-5 *Configuring DTP on SW1 and SW2*

```
SW1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SW1(config)# interface gi1/0/2
SW1(config-if)# switchport mode dynamic auto
```

```
SW2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SW2(config)# interface gi1/0/1
SW2(config-if)# switchport mode dynamic desirable
```

Example 5-6 *Verifying Dynamic Trunk Port Status*

```
SW1# show interfaces trunk
! Output omitted for brevity

Port        Mode             Encapsulation  Status        Native vlan
Gi1/0/2     auto             802.1q         trunking      1

Port        Vlans allowed on trunk
Gi1/0/2     1-4094
```

```
SW2# show interfaces trunk
! Output omitted for brevity

Port        Mode             Encapsulation  Status        Native vlan
Gi1/0/1     desirable        802.1q         trunking      1

Port        Vlans allowed on trunk
Gi1/0/1     1-4094
```

DTP Mode Off Configuration

A static trunk port attempts to establish and negotiate a trunk port with a neighbor by default.

However, the interface configuration command **switchport nonegotiate** prevents that port from forming a trunk port. Example 5-7 demonstrates the use of this command. The setting is then verified by looking at the switch port status. Negotiation of Trunk now displays as Off.

Example 5-7 *Disabling Trunk Port Negotiation*

```
SW1# show run interface gi1/0/2
Building configuration...
!
interface GigabitEthernet1/0/2
 switchport mode trunk
 switchport nonegotiate
end
```

```
SW1# show interfaces gi1/0/2 switchport | i Trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: Off
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
```

Note: As a best practice, configure both ends of a link as a fixed port type (using **switchport mode access** or **switchport mode trunk**) to remove any uncertainty about the port's operations.

EtherChannel Bundle

- Ethernet network speeds are based on powers of 10 (10 Mbps, 100 Mbps, 1 Gbps, 10 Gbps,100 Gbps).
- When a link between switches becomes saturated, how can more bandwidth be added to that link to prevent packet loss?

Multiple Links

Ideally, it would be nice to plug in a second cable and double the bandwidth between the switches. However, Spanning Tree Protocol (STP) will place one of the ports into a blocking state to prevent forwarding loops, as shown in Figure 5-2.

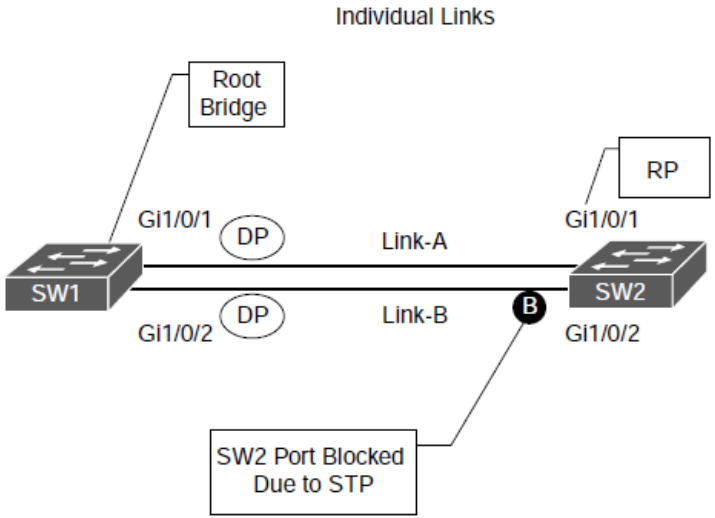

Figure 5-2 *Multiple Links with STP*

EtherChannel Components

Figure 5-3 shows some of the key components of an EtherChannel bundle between SW1 and SW2, with their Gi1/0/1 and Gi1/0/2 interfaces.

The physical links can be aggregated into a logical link called an EtherChannel bundle.

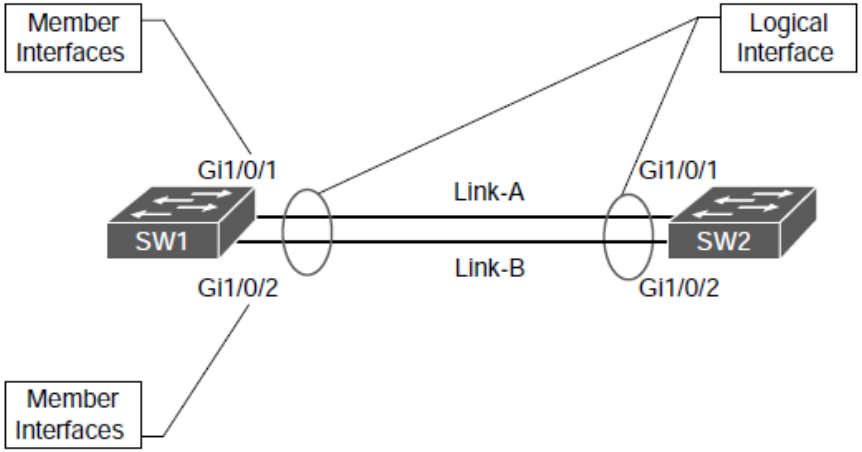

Figure 5-3 *EtherChannel Components*

EtherChannel Components (Cont.)

Aspects of EtherChannel:

- Etherchannel is defined in the IEEE 802.3AD link aggregation specification.
- STP operates on a logical link and not on a physical link.
- The logical link will have the bandwidth of any active member interfaces.
- It will load balanced across all the links.
- EtherChannels can be used for either Layer 2 (access or trunk) or Layer 3 links.

The terms EtherChannel, EtherChannel bundle, and port channel are interchanged frequently on the Catalyst platform, but other Cisco platforms only use the term port channel exclusively.

EtherChannel Link-State

- EtherChannel may be created statically or dynamically.
- Static EtherChannel does not have a health integrity check. If the physical medium degrades and keeps the line protocol in an up state, the port channel will reflect that link as viable for transferring data.
- A common scenario involves the use of intermediary devices and technologies (for example, powered network taps, IPSs, Layer 2 firewalls, DWDM) between devices. It is critical for the link state to be propagated to the other side.

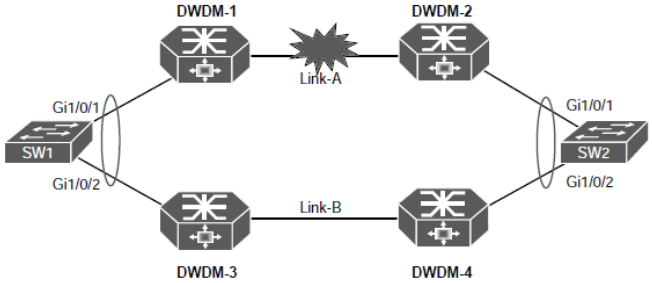

Figure 5-4 *Port-Channel Link-State Propagation and Detection*

Dynamic Link Aggregation Protocols

Two common link aggregation protocols are Link Aggregation Control Protocol (LACP) and Port Aggregation Protocol (PAgP).

- PAgP is Cisco proprietary and was developed first.
- LACP was created as an open industry standard.
- All the member links must participate in the same protocol on the local and remote switches.

PAgP Port Modes

PAgP advertises messages with the multicast MAC address 0100:0CCC:CCCC and the protocol code 0x0104. PAgP can operate in two modes:

PAgP Port Modes	Description
Auto	• The interface does not initiate an EtherChannel to be established and does not transmit PAgP packets out of it. • If an PAgP packet is received from the remote switch, this interface responds and then can establish a PAgP adjacency. • If both devices are PAgP auto, a PAgP adjacency does not form.
Desirable	• An interface tries to establish an EtherChannel and transmit PAgP packets out of it. • Active PAgP interfaces can establish a PAgP adjacency only if the remote interface is configured to auto or desirable.

LACP Port Modes

LACP advertises messages with the multicast MAC address 0180:C200:0002. LACP can operate in two modes:

LACP Port Modes	Description
Passive	• An interface does not initiate an EtherChannel to be established and does not transmit LACP packets out of it. • If an LACP packet is received from the remote switch, this interface responds and then can establish an LACP adjacency. • If both devices are LACP passive, an LACP adjacency does not form.
Active	• An interface tries to establish an EtherChannel and transmit LACP packets out of it. • Active LACP interfaces can establish an LACP adjacency only if the remote interface is configured to active or passive.

EtherChannel Configurations

It is possible to configure EtherChannels by going into the interface configuration mode for the member interfaces and assigning them to an EtherChannel ID and configuring the appropriate mode:

- **Static EtherChannel:** A static EtherChannel is configured with the interface parameter command **channel-group** *etherchannel-id* **mode on**.
- **LACP EtherChannel:** An LACP EtherChannel is configured with the interface parameter command **channel-group** *etherchannel-id* **mode** {**active** | **passive**}.
- **PAgP EtherChannel:** A PAgP EtherChannel is configured with the interface parameter command **channel-group** *etherchannel-id* **mode** {**auto** | **desirable**} [**non-silent**].
  - By default, PAgP ports operate in silent mode, which allows a port to establish an EtherChannel with a device that is not PAgP capable and rarely sends packets.
  - Using the optional **non-silent** keyword requires a port to receive PAgP packets before adding it to the EtherChannel, which is recommended.

EtherChannel Configurations (Cont.)

The following needs to be considered with EtherChannel configuration:

- Configuration settings for the EtherChannel are placed in the port-channel interface.
- Member interfaces need to be in the appropriate Layer 2 or Layer 3 (that is, no switch port) before being associated with the port channel.

Example 5-8 *Sample Port-Channel Configuration*

```
SW1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SW1(config)# interface range gi1/0/1-2
SW1(config-if-range)# channel-group 1 mode active
Creating a port-channel interface Port-channel 1
SW1(config-if-range)# interface port-channel 1
SW1(config-if)# switchport mode trunk
```

```
SW2# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SW2(config)# interface range gi1/0/1-2
SW2(config-if-range)# channel-group 1 mode passive
Creating a port-channel interface Port-channel 1
SW2(config-if-range)# interface port-channel 1
SW2(config-if)# switchport mode trunk
*13:57:05.434: %LINEPROTO-5-UPDOWN: Line protocol on Interface
  GigabitEthernet1/0/1, changed state to down
*13:57:05.446: %LINEPROTO-5-UPDOWN: Line protocol on Interface
  GigabitEthernet1/0/2, changed state to down
*13:57:12.722: %ETC-5-L3DONTBNDL2: Gi1/0/1 suspended: LACP currently not enabled
  on the remote port.
*13:57:13.072: %ETC-5-L3DONTBNDL2: Gi1/0/2 suspended: LACP currently not enabled
  on the remote port.
*13:57:24.124: %LINEPROTO-5-UPDOWN: Line protocol on Interface
  GigabitEthernet1/0/2, changed state to up
*13:57:24.160: %LINEPROTO-5-UPDOWN: Line protocol on Interface
  GigabitEthernet1/0/1, changed state to up
*13:57:25.103: %LINK-3-UPDOWN: Interface Port-channel1, changed state to up
*13:57:26.104: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel1,
  changed state to up
```

Verify Port-Channel Status

- As shown in Example 5-9, the command **show etherchannel summary** provides an overview of all the configured EtherChannels, along with the status and dynamic aggregation protocol for each one.
- When viewing the output of the **show etherchannel summary** command, the first thing that should be checked is the EtherChannel status, which is listed in the Port-channel column.
- The status should be SU, as highlighted in Example 5-9.

Example 5-9 *Viewing EtherChannel Summary Status*

```
SW1# show etherchannel summary
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator

        M - not in use, minimum links not met
        u - unsuitable for bundling
       w - waiting to be aggregated
        d - default port
        A - formed by Auto LAG

Number of channel-groups in use: 1
Number of aggregators:           1

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
1      Po1(SU)         LACP      Gi1/0/1(P)   Gi1/0/2(P)
2      Po2(SU)         PAgP      Gi1/0/3(P)   Gi1/0/4(P)
```

Note: The status codes are case sensitive, so please pay attention to the case of the field.

EtherChannel Logical Interface Status Fields

Logical EtherChannel Interface Status Fields are as follows:

- U - The EtherChannel interface is working properly.
- D - The EtherChannel interface is down.
- M - The EtherChannel interface has successfully established at least one LACP adjacency; however, the EtherChannel is configured with a minimum number of active interfaces that exceeds the number of active participating member interfaces. Traffic will not be forwarded across this port channel. The command **port-channel min-links** *min-member-interfaces* is configured on the port-channel interface.
- S - The port-channel interface is configured for Layer 2 switching.
- R - The port-channel interface is configured for Layer 3 routing.

EtherChannel Member Interface Status Fields

EtherChannel Member Interface Status Fields are as follows:

- P - The interface is actively participating and forwarding traffic for this port channel.
- H - The port-channel is configured with the maximum number of active interfaces. This interface is participating in LACP with the remote peer, but the interface is acting as a hot standby and does not forward traffic. The command **lacp max-bundle** *number-member-interfaces* is configured on the port-channel interface.
- I - The member interface has not detected any LACP activity on this interface and is treated as an individual.
- w - There is time left to receive a packet from this neighbor to ensure that it is still alive.
- s - The member interface is in a suspended state.
- r - The switch module associated with this interface has been removed from the chassis.

Port-Channel Interface Status

- The logical interface can be viewed with the command **show interface port-channel** *port-channel-id*.
- The output includes traditional interface statistics and lists the member interfaces and indicates that the bandwidth reflects the combined throughput of all active member interfaces.

Example 5-10 *Viewing Port-Channel Interface Status*

```
SW1# show interfaces port-channel 1
Port-channel1 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is 0062.ec9d.c501 (bia 0062.ec9d.c501)
  MTU 1500 bytes, BW 2000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, link type is auto, media type is
  input flow-control is off, output flow-control is unsupported
  Members in this channel: Gi1/0/1 Gi1/0/2
..
```

EtherChannel Neighbors

Example 5-11 *Viewing show etherchannel port Output*

```
SW1# show etherchannel port
! Output omitted for brevity
                Channel-group listing:
                ----------------------
! This is the header that indicates all the ports that are for the first
! EtherChannel interface. Every member link interface will be listed
Group: 1
----------
                Ports in the group:
                -------------------
! This is the first member interface for interface Po1. This interface
! is configured for LACP active
Port: Gi1/0/1
------------
Port state    = Up Mstr Assoc In-Bndl
Channel group = 1           Mode = Active          Gcchange = -
Port-channel  = Po1         GC   =   -             Pseudo port-channel = Po1
Port index    = 0           Load = 0x00            Protocol =   LACP

! This interface is configured with LACP fast packets, has a port priority
! of 32,768 and is active in the bundle.

Flags:  S - Device is sending Slow LACPDUs   F - Device is sending fast LACPDUs.
        A - Device is in active mode. P - Device is in passive mode.
```

```
Local information:
                            LACP port     Admin     Oper    Port         Port
Port      Flags   State     Priority      Key       Key     Number       State
Gi1/0/1   FA      bndl      32768         0x1       0x1     0x102        0x3F

! This interface's partner is configured with LACP fast packets, has a system-id
! of 0081.c4ff.8b00, a port priority of 32,768, and is active in the bundle
! for 0d:00h:03m:38s.

 Partner's information:
                LACP port                        Admin   Oper    Port     Port
Port      Flags Priority  Dev ID          Age    key     Key     Number   State
Gi1/0/1   FA    32768     0081.c4ff.8b00   0s     0x0     0x1     0x102    0x3F

Age of the port in the current state: 0d:00h:03m:38s

..
! This is the header that indicates all the ports that are for the second
! EtherChannel interface. Every member link interface will be listed.

Group: 2
----------
                Ports in the group:
                -------------------
! This is the first member interface for interface Po2. This interface
! is configured for PAgP desirable

Port: Gi1/0/3
------------
```

EtherChannel Neighbors (Cont.)

```
Port state    = Up Mstr In-Bndl
Channel group = 2           Mode = Desirable-Sl    Gcchange = 0
Port-channel  = Po2         GC   = 0x00020001      Pseudo port-channel = Po2
Port index    = 0           Load = 0x00            Protocol =   PAgP

! This interface is in a consistent state, has a neighbor with the
! 0081.c4ff.8b00 address and has been in the current state for 54m:45s

Flags:  S - Device is sending Slow hello. C - Device is in Consistent state.
        A - Device is in Auto mode. P - Device learns on physical port.
        d - PAgP is down.
Timers: H - Hello timer is running. Q - Quit timer is running.
        S - Switching timer is running. I - Interface timer is running.

Local information:
                            Hello    Partner  PAgP     Learning  Group
Port      Flags State   Timers  Interval Count   Priority   Method  Ifindex
Gi1/0/3   SC    U6/S7   H       30s      1         128        Any      51

Partner's information:

          Partner               Partner          Partner          Partner Group
Port      Name                  Device ID        Port       Age  Flags   Cap.
Gi1/0/3   SW2                   0081.c4ff.8b00   Gi1/0/3      1s SC      20001

Age of the port in the current state: 0d:00h:54m:45s
..
```

The output from the **show etherchannel port** command can provide too much information and slow down troubleshooting when a smaller amount of information is needed.

EtherChannel Neighbors LACP and PAgP

The command **show lacp neighbor** [**detail**] displays additional information about the LACP neighbor and includes the neighbor's system ID, system priority, and whether it is using fast or slow LACP packet intervals as part of the output.

The command **show pagp neighbor** displays additional information about the PAgP neighbor and includes the neighbor's system ID, remote port number, and whether it is using fast or slow PAgP packet intervals as part of the output.

Example 5-12 *Viewing LACP Neighbor Information*

```
SW1# show lacp neighbor
Flags:  S - Device is requesting Slow LACPDUs
        F - Device is requesting Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode

Channel group 1 neighbors

                LACP port                       Admin  Oper   Port    Port
Port      Flags Priority  Dev ID          Age   key    Key    Number  State
Gi1/0/1   SA    32768     0081.c4ff.8b00  1s    0x0    0x1    0x102   0x3D
Gi1/0/2   SA    32768     0081.c4ff.8b00  26s   0x0    0x1    0x103   0x3D
```

Example 5-13 *Viewing PAgP Neighbor Information*

```
SW1# show pagp neighbor
Flags:  S - Device is sending Slow hello. C - Device is in Consistent state.
        A - Device is in Auto mode. P - Device learns on physical port.

Channel group 2 neighbors
          Partner         Partner         Partner          Partner Group
Port      Name            Device ID       Port       Age   Flags   Cap.
Gi1/0/3   SW2             0081.c4ff.8b00  Gi1/0/3    11s   SC      20001
Gi1/0/4   SW2             0081.c4ff.8b00  Gi1/0/4    5s    SC      20001
```

Verifying EtherChannel Packets LACP and PAgP

A vital step in troubleshooting the establishment of port channels is to verify that LACP or PAgP packets are being transmitted between devices.

The first troubleshooting step that can be taken is to verify the EtherChannel counters for the appropriate protocol.

The LACP counters can be cleared with the command **clear lacp counters**. The PAgP counters can be cleared with the command **clear pagp counters**.

Example 5-14 *Viewing LACP Packet Counters*

```
SW2# show lacp counters
              LACPDUs         Marker      Marker Response    LACPDUs
Port        Sent    Recv     Sent   Recv     Sent   Recv      Pkts Err
--------------------------------------------------------------------
Channel group: 1
Gi1/0/1      23      23       0      0        0      0         0
Gi1/0/2      22       0       0      0        0      0         0
```

```
SW2# show lacp counters
              LACPDUs         Marker      Marker Response    LACPDUs
Port        Sent    Recv     Sent   Recv     Sent   Recv      Pkts Err
--------------------------------------------------------------------
Channel group: 1
Gi1/0/1      28      28       0      0        0      0         0
Gi1/0/2      27       0       0      0        0      0         0
```

Example 5-15 *Viewing PAgP Packet Counters*

```
SW1# show pagp counters
          Information         Flush         PAgP
Port      Sent     Recv     Sent     Recv     Err Pkts
-----------------------------------------------------
Channel group: 2
Gi1/0/3   31       51       0        0        0
Gi1/0/4   44       38       0        0        0
```

Advanced LACP Configuration Options

LACP provides some additional tuning that is not available with PAgP.

LACP has the following advanced settings:

- LACP Fast
- Minimum Number of Port-Channel Member Interfaces
- Maximum Number of Port-Channel Member Interfaces
- LACP System Priority
- LACP Interface Priority

LACP Fast

- The original LACP standards sent out LACP packets every 30 seconds. A link is deemed unusable if an LACP packet is not received after three intervals, 90 seconds.
- An amendment to the standards was made so that LACP packets are advertised every 1 second. This is known as *LACP fast* because a link can be identified and removed in 3 seconds compared to the 90 seconds.
- LACP fast is enabled on the member interfaces with the interface configuration command **lacp rate fast**.

Note: All the interfaces on both switches need to be configured the same (either using LACP fast or LACP slow) for the EtherChannel to successfully come up.

Example 5-16 *Configuring LACP Fast and Verifying LACP Speed State*

```
SW1(config)# interface range gi1/0/1-2
SW1(config-if-range)# lacp rate fast
```

```
SW1# show lacp internal
Flags:  S - Device is requesting Slow LACPDUs
        F - Device is requesting Fast LACPDUs
        A - Device is in Active mode       P - Device is in Passive mode

Channel group 1
                            LACP port     Admin     Oper    Port        Port
Port      Flags   State     Priority      Key       Key     Number      State
Gi1/0/1   FA      bndl      32768         0x1       0x1     0x102       0x3F
Gi1/0/2   FA      bndl      32768         0x1       0x1     0x103       0xF
```

Minimum Number of Port-Channel Member Interfaces

- An EtherChannel interface becomes active and up when only one member interface successfully forms an adjacency with a remote device.
- In some design scenarios using LACP, a minimum number of adjacencies is required before a port-channel interface becomes active.
- This option can be configured with the port-channel interface command **port-channel min-links** *min-links*.

The minimum number of port-channel member interfaces does not need to be configured on both devices to work properly. However, configuring it on both switches is recommended to accelerate troubleshooting and assist operational staff.

Example 5-17 *Configuring the Minimum Number of EtherChannel Member Interfaces*

```
SW1(config)# interface port-channel 1
SW1(config-if)# port-channel min-links 2
```

```
SW1(config-if)# interface gi1/0/1
SW1(config-if)# shutdown
10:44:46.516: %ETC-5-MINLINKS_NOTMET: Port-channel Po1 is down bundled ports (1)
   doesn't meet min-links
10:44:47.506: %LINEPROTO-5-UPDOWN: Line protocol on Interface Gigabit
   Ethernet1/0/2, changed state to down
10:44:47.508: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel1,
   changed state to down
10:44:48.499: %LINK-5-CHANGED: Interface GigabitEthernet1/0/1, changed state to
   administratively down
10:44:48.515: %LINK-3-UPDOWN: Interface Port-channel1, changed state to down
```

```
SW1# show etherchannel summary
! Output Ommitted for Brevity
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator

        M - not in use, minimum links not met
..
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------

1      Po1(SM)          LACP      Gi1/0/1(D)   Gi1/0/2(P)
```

Maximum Number of Port-Channel Member Interfaces

- An EtherChannel can be configured to have a specific maximum number of member interfaces in a port channel.
- This may be done to ensure that the active member interface count proceeds with powers of two (for example, 2, 4, 8) to accommodate load-balancing hashes.
- The maximum number of member interfaces in a port channel can be configured with the port-channel interface command **lacp max-bundle** *max-links*.

The maximum number of port-channel member interfaces needs to be configured only on the master switch for that port channel. However, configuring it on both switches is recommended to accelerate troubleshooting.

Example 5-18 *Configuring and Verifying the Maximum Links*

```
SW1(config)# interface port-channel1
SW1(config-if)# lacp max-bundle 1
11:01:11.972: %LINEPROTO-5-UPDOWN: Line protocol on Interface Gigabit
   Ethernet1/0/1, changed state to down
11:01:11.979: %LINEPROTO-5-UPDOWN: Line protocol on Interface Gigabit
   Ethernet1/0/2, changed state to down
11:01:11.982: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel1,
   changed state to down
11:01:13.850: %LINEPROTO-5-UPDOWN: Line protocol on Interface Gigabit
   Ethernet1/0/1, changed state to up
11:01:13.989: %LINEPROTO-5-UPDOWN: Line protocol on Interface Port-channel1,
   changed state to up
```

```
SW1# show etherchannel summary
! Output omitted for brevity
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator

        M - not in use, minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

        A - formed by Auto LAG
..
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
1      Pol(SU)         LACP      Gi1/0/1(P)  Gi1/0/2(H)
```

LACP System Priority

The *LACP system priority* identifies which switch is the master switch for a port channel.

- The master switch on a port channel is responsible for choosing which member interfaces are active in a port channel when there are more member interfaces than the maximum number of member interfaces associated with a port-channel interface.
- The switch with the lower system priority is preferred.
- The LACP system priority can be changed with the command **lacp system-priority** *priority*.

Example 5-19 *Viewing and Changing the LACP System Priority*

```
SW1# show lacp sys-id
32768, 0062.ec9d.c500
```

```
SW1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
SW1(config)# lacp system-priority 1
```

```
SW1# show lacp sys-id
1, 0062.ec9d.c500
```

LACP Interface Priority

LACP interface priority enables the master switch to choose which member interfaces are active in a port channel when there are more member interfaces than the maximum number of member interfaces for a port channel.

- A port with a lower port priority is preferred.
- The interface configuration command **lacp port-priority** *priority* sets the interface priority.

Example 5-20 *Changing the LACP Port Priority*

```
SW1# show etherchannel summary | b Group
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
1      Po1(SU)         LACP      Gi1/0/1(P)  Gi1/0/2(H)
```

```
SW1(config)# interface gi1/0/2
SW1(config-if)# lacp port-priority 1
SW1# show etherchannel summary | b Group
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
1      Po1(SU)         LACP      Gi1/0/1(H)  Gi1/0/2(P)
```

Troubleshooting EtherChannel Bundles

A port channel is a logical interface, so all the member interfaces must have the same characteristics. If they do not, problems will occur.

- Generally as a rule, when configuring port channels on a switch, place each member interface in the appropriate switch port type (Layer 2 or Layer 3) and then associate the interfaces to a port channel.

- All other port-channel configuration is done via the port-channel interface.

Troubleshooting EtherChannel Bundles

The following configuration settings must match on the member interfaces:

Match Settings	
Port type	Duplex
Port mode	MTU
Native VLAN	Load interval
Allowed VLAN	Storm control
Speed	

In addition to the matched configuration settings, check the following when troubleshooting the establishment of an EtherChannel bundle:

- Ensure that a member link is between only two devices.
- Ensure that the member ports are all active.
- Ensure that both end links are statically set to *on* or that either LACP is enabled with at least one side set to *active* or PAgP is enabled with at least one side set to *desirable*.
- Ensure that all member interface ports are consistently configured (except for LACP port priority).
- Verify the LACP or PAgP packet transmission and receipt on both devices.

Load Balancing Traffic with EtherChannel Bundles

- Traffic that flows across a port-channel interface is not forwarded out member links on a round-robin basis per packet.
- Instead, a hash is calculated, and packets are consistently forwarded across a link based on that hash, which runs on the various packet header fields.
- The load-balancing hash is a system wide configuration that uses the global command **port-channel load-balance** *hash*.

Port-Channel Load Balancing Hash Option

The global command **port-channel load-balance** *hash.* The *hash* option has the following keyword choices:

- **dst-ip:** Destination IP address
- **dst-mac:** Destination MAC address
- **dst-mixed-ip-port:** Destination IP address and destination TCP/UDP port
- **dst-port:** Destination TCP/UDP port
- **src-dst-ip:** Source and destination IP addresses
- **src-dest-ip-only:** Source and destination IP addresses only
- **src-dst-mac:** Source and destination MAC addresses
- **src-dst-mixed-ip-port:** Source and destination IP addresses and source and destination TCP/UDP ports
- **src-dst-port:** Source and destination TCP/UDP ports only
- **src-ip:** Source IP address
- **src-mac:** Source MAC address
- **src-mixed-ip-port:** Source IP address and source TCP/UDP port
- **src-port:** Source TCP/UDP port

Viewing Port-Channel Hash Algorithm

If the links are unevenly distributed, changing the hash value may provide a different distribution ratio across member links.

- For example, if a port channel is established with a router, using a MAC address as part of the hash could impact the traffic flow as the router's MAC address does not change (as the MAC address for the source or destination will always be the router's MAC address).
- A better choice would be to use the source/destination IP address or base the hash on TCP/UDP session ports.

The command **show etherchannel load-balance** displays how a switch will load balance network traffic based on its type: non-IP, IPv4, or IPv6.

Example 5-21 *Viewing the Port-Channel Hash Algorithm*

```
SW1# show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
        src-dst-mixed-ip-port

EtherChannel Load-Balancing Addresses Used Per-Protocol:
Non-IP: Source XOR Destination MAC address
  IPv4: Source XOR Destination IP address and TCP/UDP (layer-4) port number
  IPv6: Source XOR Destination IP address and TCP/UDP (layer-4) port number
```

A hash is a binary function, so links should be in powers of two (for example, 2, 4, 8), to be consistent.

- A three-port EtherChannel will not load balance as effectively as a two- or four-port EtherChannel.
- The best was to view the load of each member link is with the command **show etherchannel port**.

Prepare for the Exam

Prepare for the Exam

Key Topics for Chapter 5

Description	
VLAN Trunking Protocol (VTP)	Minimum number of port-channel member interfaces
VTP revision reset	Maximum number of port-channel member interfaces
Dynamic Trunking Protocol (DTP)	LACP system priority
Disabling DTP	LACP interface priority
PAgP port modes	Troubleshooting EtherChannel Bundles
LACP port modes	Load balancing traffic with EtherChannel bundles
EtherChannel configuration	

Prepare for the Exam

Key Terms for Chapter 5

Terms
Dynamic Trunking Protocol (DTP)
EtherChannel bundle
member links
LACP interface priority
LACP system priority
load-balancing hash
VLAN Trunking Protocol (VTP)

Command Reference for Chapter 5

Task	Command Syntax
Configure the VTP version	**vtp version** {**1** | **2** | **3**}
Configure the VTP domain name	**vtp domain** *domain-name*
Configure the VTP mode for a switch	**vtp mode** { **server** | **client** | **transparent** | **none**} (required for the first VTP v3 server) **vtp primary**
Display the STP root bridge and cost	**switchport mode dynamic desirable**
Configure a switch port to actively attempt to establish a trunk link	**switchport mode dynamic auto**
Configure the member ports for a static EtherChannel	**channel-group** *etherchannel-id* **mode on**

Command Reference for Chapter 5 (Cont.)

Task	Command Syntax
Configure the member ports for a LACP EtherChannel	**channel-group** *etherchannel-id* **mode** {**active** | **passive**}
Configure the member ports for a PAgP EtherChannel	**channel-group** *etherchannel-id* **mode** {**auto** | **desirable**} [**non-silent**]
Configure the LACP packet rate	**lacp rate** {**fast** | **slow**}
Configure the minimum number of member links for the LACP EtherChannel to become active	**port-channel min-links** *min-links*
Configure the maximum number of member links in an LACP EtherChannel	**lacp max-bundle** *max-links*

Command Reference for Chapter 5 (Cont.)

Task	Command Syntax
Configure a switch's LACP system priority	**lacp system-priority** *priority*
Configure a switch's LACP port priority	**lacp port-priority** *priority*
Configure the EtherChannel load-balancing hash algorithm	**port-channel load-balance** *hash*
Display the contents of all current access lists	**show access-list** [*access-list-number* | *access-list-name*}
Display the VTP system settings	**show vtp status**
Display the switch port DTP settings, native VLANs, and allowed VLANs	**show interface** [*interface-id*] **trunk**
Display a brief summary update on EtherChannel interfaces	**show etherchannel summary**

Command Reference for Chapter 5 (Cont.)

Task	Command Syntax
Display detailed information for the local EtherChannel interfaces and their remote peers	**show interface port-channel**
Display information about LACP neighbors	**show lacp neighbor** [**detail**]
Display the local LACP system identifier and priority	**show lacp** *system-id*
Display the LACP counters for configure interfaces	**show lacp counters**
Display information about PAgP neighbors	**show pagp neighbor**
Display the PAgP counters for configured interfaces	**show pagp counters**
Display the algorithm for load balancing network traffic based on the traffic type	**show etherchannel load-balance**


CISCO