

CCNP-RS SWITCH v2.0

CCNP-RS SWITCH Timetable

Day	AM	Lunch	PM
1	VLANs, Trunking and VTP	--	<i>VLANs, Trunk, VTP Labs</i>
2	STP	--	<i>STP, MSTP Labs</i>
3	Inter Vlan Routing, DHCP and HA	--	<i>Routing, DHCP and HA Labs</i>
4	Security and Monitoring	--	<i>Security and Monitoring Labs</i>
5	Intro MPLS Forwarding (Extra)	--	<i>MPLS L3VPN and L2VPN Labs</i>

Chapter 1

Implementing VLANs in Campus Networks

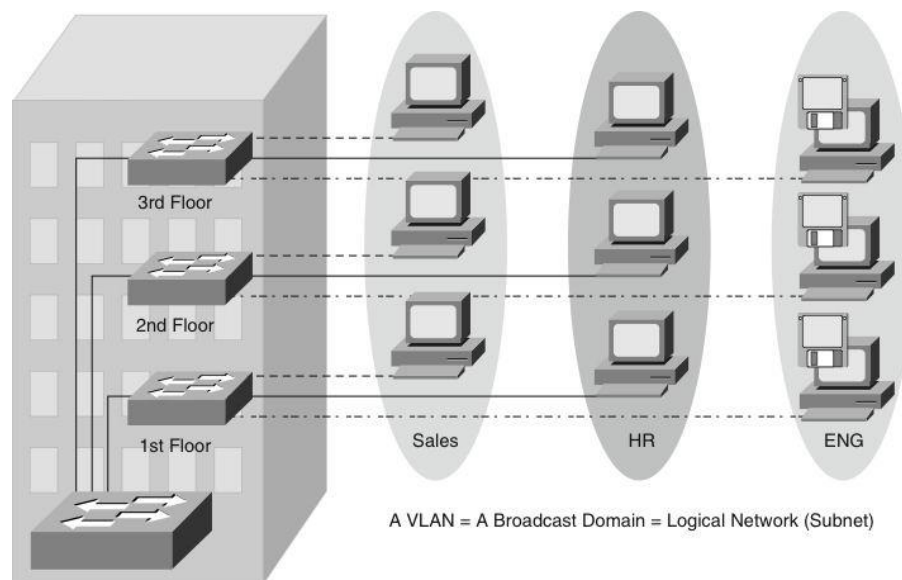
CCNP-RS SWITCH

Chapter 1 Objectives

- Design and plan VLANs, trunks, and addressing to meet business requirements, technical requirements, and constraints.
- Configure VLANs and VLAN trunks in the campus network to support business and technical requirements.
- Configure VTP in the campus network to support business and technical requirements.
- Describe private VLANs and configure private VLANs in the campus network to support business and technical requirements.
- Configure and verify an EtherChannel in a Layer 2 topology that contains bridging loops.


Implementing VLAN Technologies in a Campus Network

Virtual Local Area Network (VLAN)




- A VLAN is a logical group of end devices.
- Broadcasts are contained within VLANs.
- Modern design has 1 VLAN = 1 IP subnet.
- Trunks connect switches so as to transport multiple VLANs.
- Layer 3 devices interconnect VLANs.

End-to-End VLANs



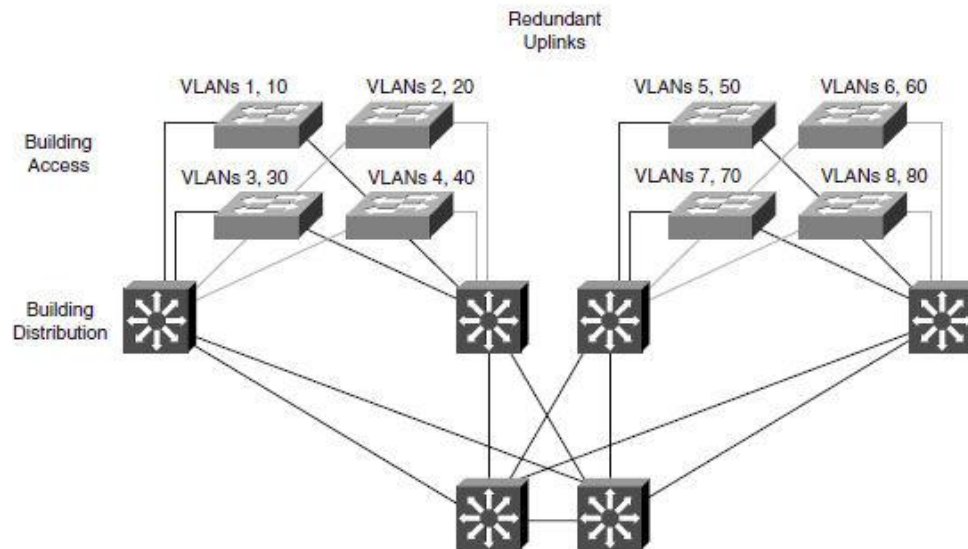
- Each VLAN is distributed geographically throughout the network.
- Users are grouped into each VLAN regardless of the physical location, theoretically easing network management.
- As a user moves throughout a campus, the VLAN membership for that user remains the same.
- Switches are configured for VTP server or client mode.

Local VLANs



- Create local VLANs with physical boundaries in mind rather than job functions of the users.
- Local VLANs exist between the access and distribution layers.
- Traffic from a local VLAN is routed at the distribution and core levels.
- Switches are configured in VTP transparent mode.
- Spanning tree is used only to prevent inadvertent loops in the wiring closet.
- One to three VLANs per access layer switch recommended.

VLANs in Enterprise Campus Design



- VLANs used at the access layer should extend no further than their associated distribution switch.
- Traffic is routed from the local VLAN as it is passed from the distribution layer into the core.
- Blocks can contain one to three VLANs each.
- STP is limited to access and distribution switches.
- DHCP is used to assign IP addresses to users.

Best Practices for VLAN Design

- One to three VLANs per access module and limit those VLANs to a couple of access switches and the distribution switches.
- Avoid using **VLAN 1** as the "blackhole" for all unused ports. Use a dedicated VLAN separate from VLAN 1 to assign all the unused ports.
- Separate the **voice VLANs**, **data VLANs**, the **management VLAN**, the native VLAN, blackhole VLANs, and the default VLAN (VLAN 1).
- Avoid VTP when using local VLANs; use manually allowed VLANs on trunks.
- For trunk ports, **turn off** Dynamic Trunking Protocol (**DTP**) and configure trunking. Use IEEE 802.1Q rather than ISL because it has better support for QoS and is a standard protocol.
- Manually configure access ports that are not specifically intended for a trunk link.
- Prevent all data traffic from VLAN 1; only permit control protocols to run on **VLAN 1 (DTP, VTP, STP BPDUs, PAgP, LACP, CDP, etc.)**.
- Avoid using **Telnet** because of security risks; enable **SSH** support on management VLANs.

VLAN Support on Catalyst Switches

Catalyst Switch	Max VLANs	VLAN ID Range
2940	4	1 - 1005
2950/2955	250	1 - 4094
2960	255	1 - 4094
2970/3550/3560/3750	1055	1 - 4094
2848G/2980G/4000/4500	4094	1 - 4094
6500	4094	1 - 4094

VLAN Ranges on Catalyst Switches

VLAN Range	Range	Usage	Popagated via VTP?
0, 4095	Reserved	For system use only. You cannot see or use these.	n/a
1	Normal	Cisco default. You can use this VLAN, but you cannot delete it.	Yes
2 – 1001	Normal	For Ethernet VLANs. You can create, use, and delete these.	Yes
1002 – 1005	Normal	Cisco defaults for FDDI and Token Ring. You cannot delete these.	Yes
1006 – 1024	Reserved	For system use only. You cannot see or use these.	n/a
1025 - 4094	Reserved	For Ethernet VLANs only.	VTP v 3 only. Not supported in VTP v1 or v2. Requires VTP transparent mode for configuration.

Configuration: Create a VLAN

- To create a new VLAN in global configuration mode.

```
Switch(config)# vlan vlan-id
```

- *vlan-id* is 2-1001 or 1025-4094

Configuration: Name a VLAN

- To name a VLAN in VLAN configuration mode.
`Switch(config-vlan)# name vlan-name`
- *vlan-name* is a descriptor for the VLAN.
- Naming a VLAN is optional.

Example: Creating and Naming a VLAN

- Enter global configuration mode:
Switch# **configure terminal**
- Create a new VLAN with a particular ID number:
Switch(config)# **vlan** *vlan-id*
- (Optional.) Name the VLAN:
Switch(config-vlan)# **name** *vlan-name*

```
Switch# configure terminal  
Switch(config)# vlan 5  
Switch(config-vlan)# name Engineering  
Switch(config-vlan)# exit
```

Configuration: Disable Trunk Negotiation on a Port

- To disable trunk negotiation on a switch port.

```
Switch(config-if) # switchport mode access
```

- This command is optional but is recommended for security purposes. An access port does not need to negotiate trunk formation.

Configuration: Macro for Access Port

- To configure an optional macro for switch access ports.

```
Switch(config-if) # switchport host
```

- This command optimizes a Layer 2 port for a host connection.
- This macro sets the port mode to access, enables spanning-tree portfast, and disables EtherChannel.

Configuration: Assign Port to VLAN

- To assign a port to a VLAN in interface configuration mode.

```
Switch(config-if)# switchport access vlan vlan-id
```

- *vlan-id* is a previously created VLAN.

Example: Assigning a Port to a VLAN

- Enter interface configuration mode:
Switch(config)# **interface** *interface-id*
- Configure a description for the device(s) connected to the port:
Switch(config-if)# **description** *string*
- Configure access port macro:
Switch(config-if)# **switchport host**
- Assign port to VLAN:
Switch(config-if)# **switchport access vlan** *vlan-id*
- Enable the interface:
Switch(config-if)# **no shutdown**
- Return to Privileged EXEC mode
Switch(config-if)# **end**

```
Switch(config)# interface FastEthernet 5/6
Switch(config-if)# description PC A
Switch(config-if)# switchport host
switchport mode will be set to access
spanning-tree portfast will be enabled
channel group will be disabled
Switch(config-if)# switchport access vlan 200
Switch(config-if)# no shutdown
Switch(config-if)# end
```

Verification: VLAN Configuration

- The `show vlan` command and its derivatives are the most useful commands for displaying information related to VLANs. The following two forms have the same output.

```
Switch# show vlan id 3
```

```
VLAN Name                Status    Ports
-----
3      VLAN0003                active    Fa0/1

VLAN Type  SAID       MTU   Parent RingNo BridgeNo  Stp  BrdgMode Trans1 Trans2
-----
3      enet  100003    1500  -     -        -    -        -        0      0
```

```
Switch# show vlan name VLAN0003
```

```
VLAN Name                Status    Ports
-----
3      VLAN0003                active    Fa0/1

VLAN Type  SAID       MTU   Parent RingNo BridgeNo  Stp  BrdgMode Trans1 Trans2
-----
3      enet  100003    1500  -     -        -    -        -        0      0
```

Verification: Interface Configuration

- The `show running-config` command has an `interface` keyword option to allow for interface-specific output.

```
Switch# show running-config interface FastEthernet 5/6  
Building configuration...  
!  
Current configuration :33 bytes  
interface FastEthernet 5/6  
switchport access vlan 200  
switchport mode access  
switchport host  
end
```

Verification: Switch Port Configuration

- One of the most useful commands for showing VLAN configuration information specific to a switch port is the **show interfaces *interface_id* switchport** command.

```
Switch# show interfaces f0/18 switchport  
Name: Fa0/18  
Switchport: Enabled  
Administrative Mode: static access  
Operational Mode: down  
Administrative Trunking Encapsulation: dot1q  
Negotiation of Trunking: Off  
Access Mode VLAN: 20 (VLAN0020)  
Trunking Native Mode VLAN: 1 (default)  
Administrative Native VLAN tagging: enabled  
Voice VLAN: 150 (VLAN0150)  
<output omitted>  
Operational private-vlan: none  
Trunking VLANs Enabled: ALL  
Pruning VLANs Enabled: 2-1001  
Capture Mode Disabled  
Capture VLANs Allowed: ALL
```

Verification: MAC Address Information

- You can view MAC address information specific to an interface and an associated VLAN.

```
Switch# show mac-address-table interface GigabitEthernet 0/1 vlan 1
```

```
Mac Address Table
```

```
-----
```


Vlan	Mac Address	Type	Ports
----	-----	----	-----
1	0008.2199.2bc1	DYNAMIC	Gi0/1

```
Total Mac Addresses for this criterion: 1
```

Implementing Trunking in a Campus Network

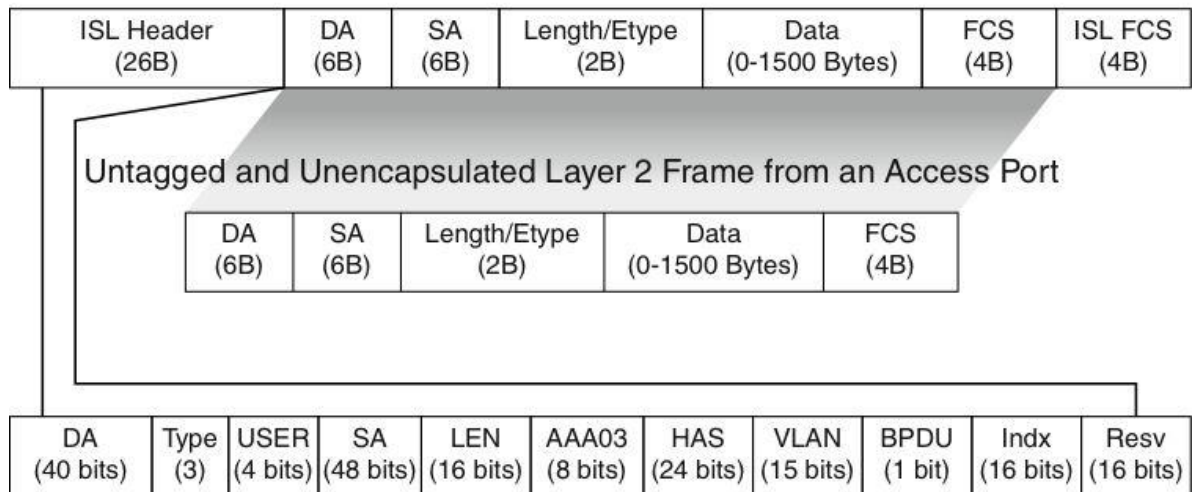
VLAN Trunking

- Trunks carry the traffic for multiple VLANs across a single physical link (multiplexing). Trunking is used to extend Layer 2 operations across an entire network.
- The host on the left in VLAN 2 can communicate with the host on the right in VLAN 2 via the trunk link; over the same trunk link, the hosts on VLAN 1 can communicate simultaneously.



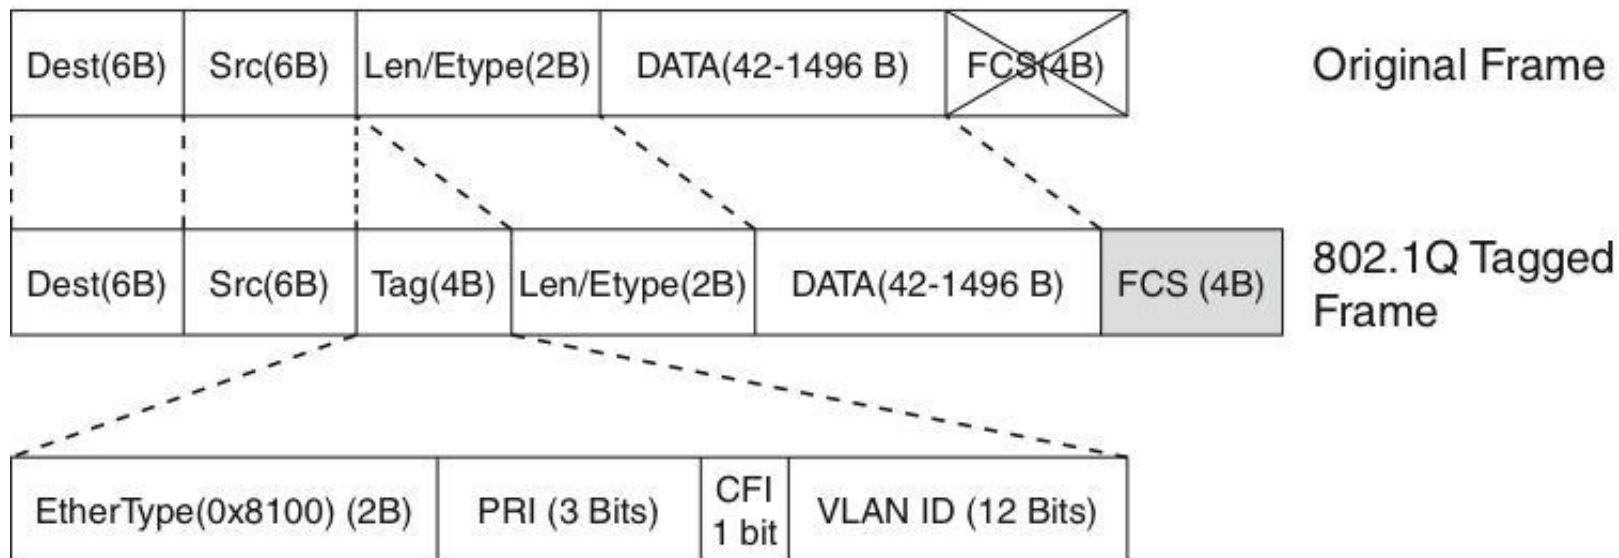
VLAN Trunking with Inter-Switch Link (ISL)

ISL Encapsulated Layer 2 Frame from an ISL Trunk Port




- ISL is Cisco-proprietary trunking protocol.
- ISL is nearly obsolete.
- ISL encapsulates Ethernet frames, adding 30 bytes of overhead.
- ISL is supported on non-access-layer Cisco switches.

VLAN Trunking with IEEE 802.1Q



- 802.1Q is a widely supported industry-standard protocol.
- IEEE 802.1Q has smaller frame overhead than ISL. 802.1Q overhead is 4 bytes.
- 802.1Q has the 802.1p field for QoS support.

Native VLAN with IEEE 802.1Q



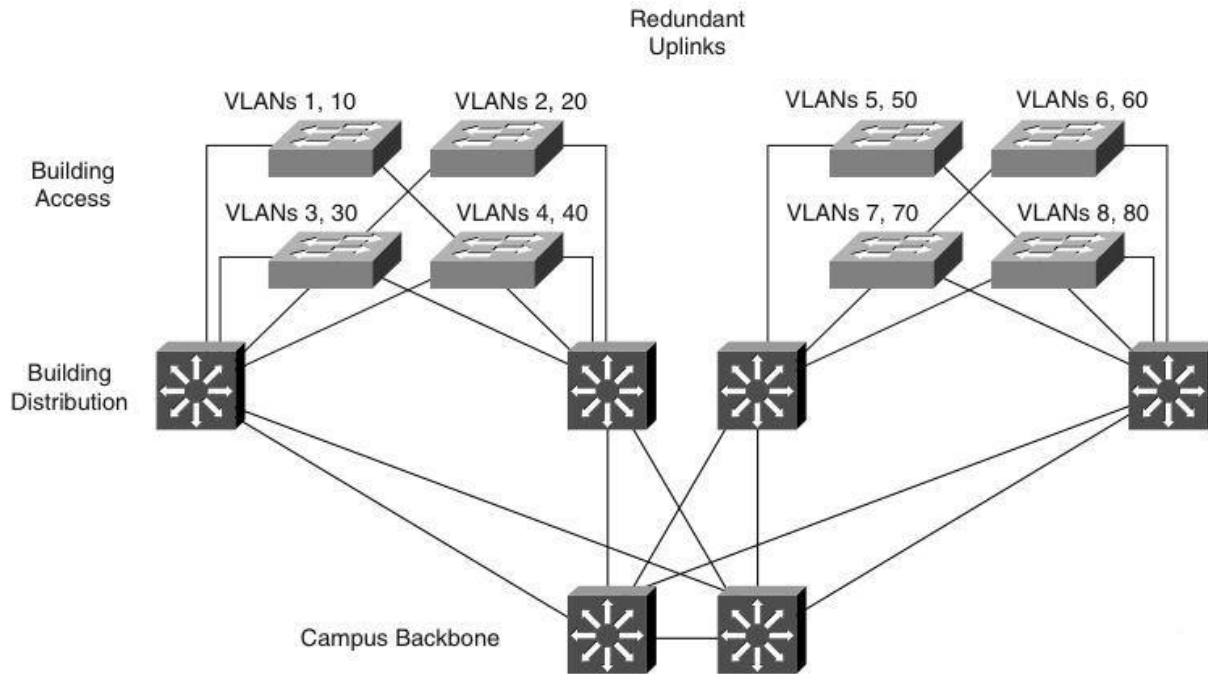
- The 802.1Q standard specifies how the switch should handle untagged frames sent or received on an 802.1Q trunk port.
- An 802.1Q trunk port is assigned a default PVID, which is associated with all untagged traffic on the port. All traffic with a null VLAN ID is assumed to belong to the port default PVID. A packet with a VLAN ID equal to the outgoing port default PVID is sent untagged. All other traffic is sent with a VLAN tag.
- Proactively configuring both ends of an 802.1Q trunk link with a native VLAN distinct from all other VLANs is recommended.

Dynamic Trunking Protocol (DTP)

	Dynamic Auto	Dynamic Desirable	Trunk	Access
Dynamic Auto	Access	Trunk	Trunk	Access
Dynamic Desirable	Trunk	Trunk	Trunk	Access
Trunk	Trunk	Trunk	Trunk	Limited connectivity
Access	Access	Access	Limited connectivity	Access

- **Access** - Puts the interface into permanent non-trunking mode and negotiates to convert the link into a non-trunk link. The interface becomes a non-trunk interface even if the neighboring interface does not agree to the change.
- **Trunk** - Puts the interface into permanent trunking mode and negotiates to convert the link into a trunk link. The interface becomes a trunk interface even if the neighboring interface does not agree to the change.
- **Nonegotiate** - Puts the interface into permanent trunking mode but prevents the interface from generating DTP frames. You must configure the neighboring interface manually as a trunk interface to establish a trunk link. Use this mode when connecting to a device that does not support DTP.
- **Dynamic desirable** - Makes the interface actively attempt to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk, desirable, or auto mode.
- **Dynamic auto** - Makes the interface willing to convert the link to a trunk link. The interface becomes a trunk interface if the neighboring interface is set to trunk or desirable mode. This is the default mode for all Ethernet interfaces in Cisco IOS.

Design with VLAN Trunks



- Trunks interconnect access layer switches.
- Trunks connect access layer switches to distribution layer switches.
- Layer 3 links interconnect core and distribution layer switches.
- Access layer switches are configured in a spanning-tree, loop-free, V-shaped topology. If one distribution link fails, HSRP or VRRP provide an alternative default gateway.
- *Recommended:* turn off DTP and manually prune VLANs on trunks.

Configuring an Interface for Trunking

- Select the encapsulation type:

```
Switch(config-if) # switchport trunk encapsulation {isl | dot1q | negotiate}
```

- Configure the interface as a Layer 2 trunk:

```
Switch(config-if) # switchport mode {dynamic {auto | desirable} | trunk}
```

- Specify the native VLAN:

```
Switch(config-if) # switchport trunk native vlan vlan-id
```

- Configure the allowable VLANs for this trunk:

```
Switch(config-if) # switchport trunk allowed vlan {add | except | all | remove} vlan-id[,vlan-id[,vlan-id[,...]]]
```

```
Switch(config)# interface FastEthernet 5/8
Switch(config-if) # switchport trunk encapsulation dot1q
Switch(config-if) # switchport mode trunk
Switch(config-if) # switchport nonegotiate ← optional
Switch(config-if) # switchport trunk allowed vlan 1-100
Switch(config-if) # no shutdown
Switch(config-if) # end
```

Verifying Trunk Configuration

```
Switch# show running-config interface f5/8
```

```
Building configuration...
Current configuration:
!
interface FastEthernet5/8
switchport mode dynamic desirable
switchport trunk encapsulation dot1q
end
```

```
Switch# show interfaces f5/8 switchport
```

```
Name: Fa5/8
Switchport: Enabled
Administrative Mode: dynamic desirable
Operational Mode: trunk
Administrative Trunking Encapsulation: negotiate
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: Enabled
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: ALL
Pruning VLANs Enabled: 2-1001
```

```
Switch# show interfaces f5/8 trunk
```

Port	Mode	Encapsulation	Status	Native vlan
Fa5/8	desirable	n-802.1q	trunking	1

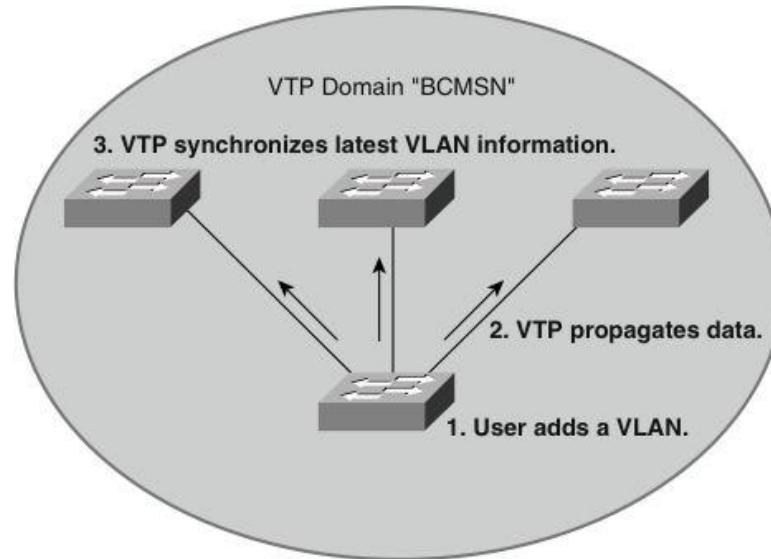
Port	Vlans allowed on trunk
Fa5/8	1-1005

Troubleshooting Trunk Links

- Ensure that the Layer 2 interface mode configured on both ends of the link is valid. The trunk mode should be **trunk** or **desirable** for at least one side of the trunk.
- Ensure that the trunk encapsulation type configured on both ends of the link is valid and compatible.
- On IEEE 802.1Q trunks, make sure the native VLAN is the same on both ends of the trunk.
- When using DTP, ensure that both ends of the link are in the same VTP domain.

VLAN Trunking Protocol

VLAN Trunking Protocol (VTP)



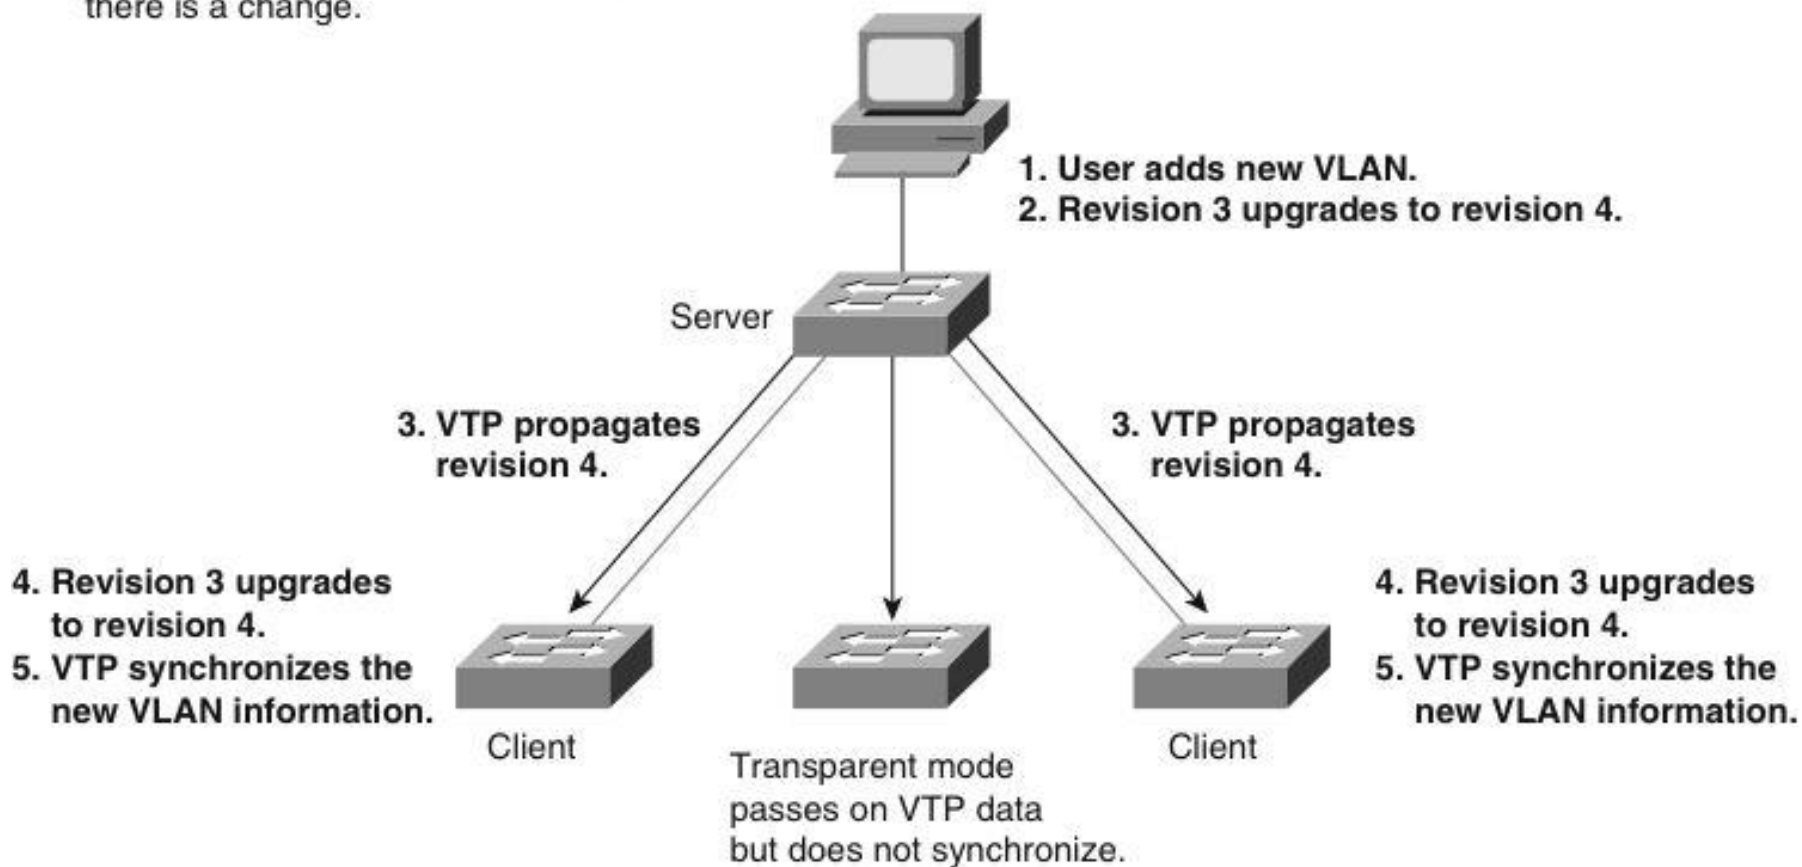
- VTP is a Cisco-proprietary protocol that automates the propagation of VLAN information between switches via trunk links. This minimizes misconfigurations and configuration inconsistencies.
- VTP does not configure switch ports for VLAN membership.
- Three types of VTP messages are sent via Layer 2 multicast on VLAN 1.
- VTP *domains* define sets of interconnected switches sharing the same VTP configuration.

VTP Modes

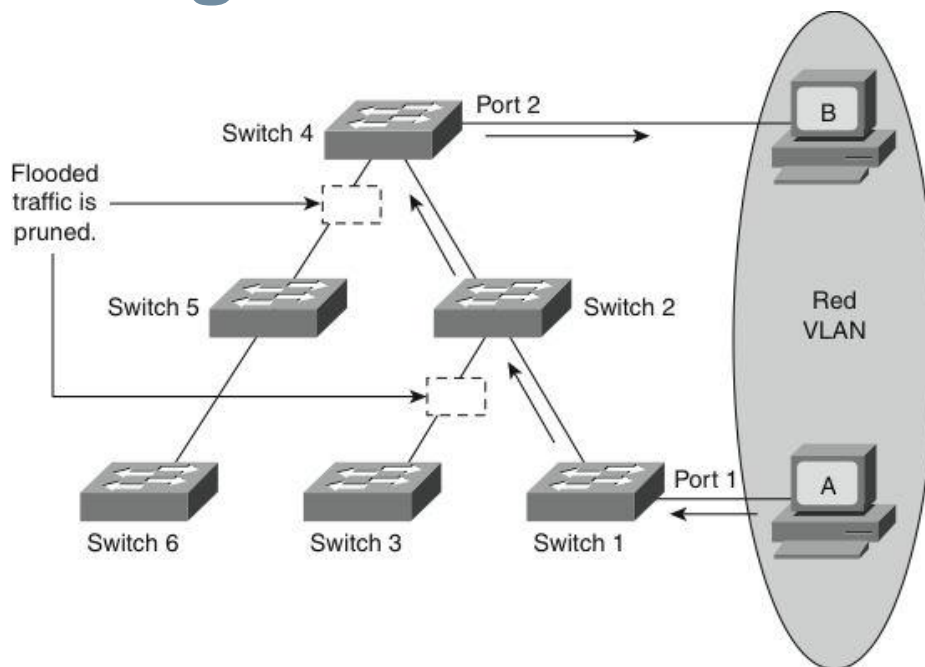
Mode	Description
Client	<ul style="list-style-type: none"> • Cannot create, change, or delete VLANs on command-line interface (CLI). • Forwards advertisements to other switches. • Synchronizes VLAN configuration with latest information received from other switches in the management domain. • Does not save VLAN configuration in nonvolatile RAM (NVRAM).
Server	<ul style="list-style-type: none"> • Can create, modify, and delete VLANs. • Sends and forwards advertisements to other switches. • Synchronizes VLAN configuration with latest information received from other switches in the management domain. • Saves VLAN configuration in NVRAM.
Transparent	<ul style="list-style-type: none"> • Can create, modify, and delete VLANs only on the local switch. • Forwards VTP advertisements received from other switches in the same management domain. • Does not synchronize its VLAN configuration with information received from other switches in the management domain. • Saves VLAN configuration in NVRAM.

VTP Operation

VTP advertisements are sent as multicast frames.
 VTP servers and clients are synchronized to the latest revision number.
 VTP advertisements are sent every 5 minutes or when there is a change.



VTP Pruning



- VTP pruning prevents flooded traffic from propagating to switches that do not have members in specific VLANs.
- VTP pruning uses VLAN advertisements to determine when a trunk connection is flooding traffic needlessly. Switches 1 and 4 in the figure support ports statically configured in the Red VLAN.
- The broadcast traffic from Station A is not forwarded to Switches 3, 5, and 6 because traffic for the Red VLAN has been pruned on the links indicated on Switches 2 and 4.

VTP Versions

- Three VTP versions: V1, V2, V3.
- Versions are not interoperable (e.g., V2 supports token ring VLANs but V1 does not).
- Unrecognized *Type-Length-Value (TLV)* configuration changes are propagated by V2 servers and clients and these unrecognized TLVs can be stored in NVRAM.
- V1 transparent switches inspect VTP messages for the domain name and version and forward a message only if the version and domain name match. V2 transparent switches forward VTP messages in transparent mode without checking versions.
- V2 performs VLAN consistency checks (VLAN names and values) only when you enter new information through the CLI or via SNMP. V2 does not perform checks when new information is obtained from a VTP message or when information is read from NVRAM. If the MD5 hash on a received VTP message is correct, V2 accepts the VTP message information.

VTP Message Types

- Summary Advertisements
- Subset Advertisements
- Advertisement Requests

VTP Summary Advertisements

Summary Advertisement			
Version	Code	Followers	MgmtD Len
Management Domain Name (Zero-Padded to 32 Bytes)			
Configuration Revision Number			
Updater Identity			
Update Timestamp (12 Bytes)			
MD5 Digest (16 Bytes)			

- By default, Catalyst switches issue summary advertisements in 5-minute increments. Summary advertisements inform adjacent switches of the current VTP domain name and the configuration revision number.
- When the switch receives a summary advertisement packet, the switch compares the VTP domain name to its own VTP domain name. If the name is different, the switch ignores the packet. If the name is the same, the switch then compares the configuration revision to its own revision. If its own configuration revision is higher or equal, the packet is ignored. If it is lower, an advertisement request is sent.

VTP Subset Advertisements

Subset Advertisements			
Version	Code	Seq-Number	Domain Name Length
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number			
VLAN-info Field 1			
:			
VLAN-info Field N			

- When you add, delete, or change a VLAN, the VTP server where the changes are made increments the configuration revision and issues a summary advertisement. One or several subset advertisements follow the summary advertisement.
- A subset advertisement contains a list of VLAN information. If there are several VLANs, more than one subset advertisement can be required to advertise all the VLANs.

VTP Subset Advertisements



- When you add, delete, or change a VLAN, the VTP server where the changes are made increments the configuration revision and issues a summary advertisement. One or several subset advertisements follow the summary advertisement.
- A subset advertisement contains a list of VLAN information. If there are several VLANs, more than one subset advertisement can be required to advertise all the VLANs.

VTP Advertisement Requests

- A switch issues a VTP advertisement request in these situations:
 - The switch has been reset.
 - The VTP domain name has been changed.
 - The switch has received a VTP summary advertisement with a higher configuration revision than its own.
- Upon receipt of an advertisement request, a VTP device sends a summary advertisement.
- One or more subset advertisements follow the summary advertisement.



VTP Authentication

- VTP domains can be secured by using the VTP password feature. It is important to make sure that all the switches in the VTP domain have the same password and domain name; otherwise, a switch will not become a member of the VTP domain. Cisco switches use MD5 to encode passwords in 16-byte words. These passwords propagate inside VTP summary advertisements. In VTP, passwords are case-sensitive and can be 8 to 64 characters in length. The use of VTP authentication is a recommended practice.
- By default, a Catalyst switch does not have a VTP password. The switch does not automatically set the password parameter, unlike other parameters that are set automatically when a VTP advertisement is received.

Configuring VTP

- **Step 1.** Enter global configuration mode:

```
Switch# configure terminal
```
- **Step 2.** Configure the VTP mode as server:

```
Switch(config)# vtp mode server
```
- **Step 3.** Configure the domain name:

```
Switch(config)# vtp domain domain_name
```
- **Step 4.** (Optional.) Enable VTP version 2:

```
Switch(config)# vtp version 2
```
- **Step 5.** (Optional.) Specify a VTP password:

```
Switch(config)# vtp password password_string
```
- **Step 6.** (Optional.) Enable VTP pruning in the management domain:

```
Switch(config)# vtp pruning
```

VTP Configuration Example

- This example creates a VTP server with domain name `Modular_Form`, password `genus`, and pruning enabled.

```
Switch# configure terminal
Switch(config)# vtp mode server
Setting device to VTP SERVER mode.
Switch(config)# vtp domain Modular_Form
Switch(config)# vtp password genus
Switch(config)# vtp pruning
Switch(config)# end
```

Verifying VTP Configuration (1)

- The most useful command for verifying VTP configuration is the **show vtp status** command. The output displayed includes the VTP version, the VTP configuration revision number, the number of VLANs supported locally, the VTP operating mode, the VTP domain name, and the VTP pruning mode.

```

Switch# show vtp status
VTP Version : 2
Configuration Revision : 247
Maximum VLANs supported locally : 1005
Number of existing VLANs : 33
VTP Operating Mode : Server
VTP Domain Name : Modular_Form
VTP Pruning Mode : Enabled
VTP V2 Mode : Disabled
VTP Traps Generation : Disabled
MD5 digest : 0x45 0x52 0xB6 0xFD 0x63 0xC8 0x49 0x80
Configuration last modified by 0.0.0.0 at 8-12-99 15:04:4
  
```


Verifying VTP Configuration (2)

- Use the **show vtp counters** command to display statistics about VTP operation. If there are any problems regarding the VTP operation, this command helps look for VTP message type updates.

```
Switch# show vtp counters
VTP statistics:
Summary advertisements received : 7
Subset advertisements received : 5
Request advertisements received : 0
Summary advertisements transmitted : 997
Subset advertisements transmitted : 13
Request advertisements transmitted : 3
Number of config revision errors : 0
Number of config digest errors : 0
Number of V1 summary errors : 0

VTP pruning statistics:

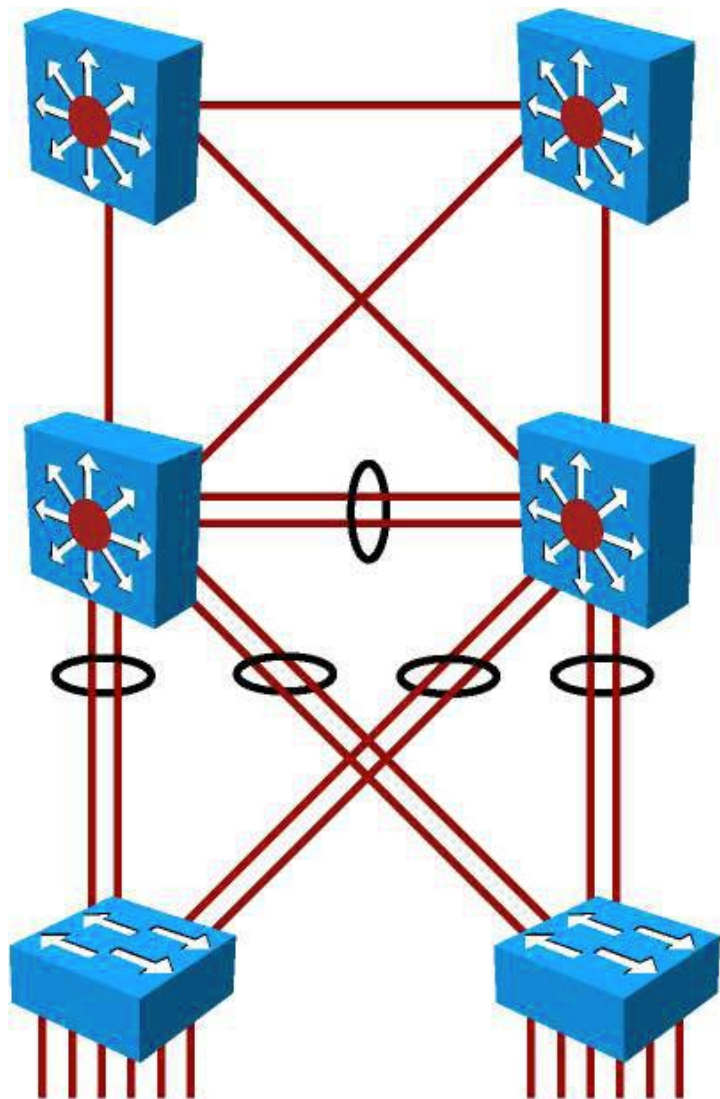
Trunk  Join Transmitted  Join Received  Summary advts received from non-pruning-capable
device
-----  -
Fa5/8    43071             42766         5
```

VTP Troubleshooting

- Check that switches are interconnected by active trunk links.
- Check that the trunking protocol matches on opposite ends of a trunk link.
- Check VTP domain name (case-sensitive) and password.
- Check the VTP mode of the switches.
- Check the VTP versions of the switches.

Configuring Link Aggregation with Etherchannel

EtherChannel Technology




334P_534

- Up to 8 physical links can be bundled into a single logical EtherChannel link.
- Usually EtherChannel is used for trunk links.
- Configuration applied to port channel interface affects all physical interfaces assigned to the port channel.
- Load balancing takes place between the physical links in an EtherChannel.
- EtherChannels can be L2 or L3 interfaces.

EtherChannel Management Protocols

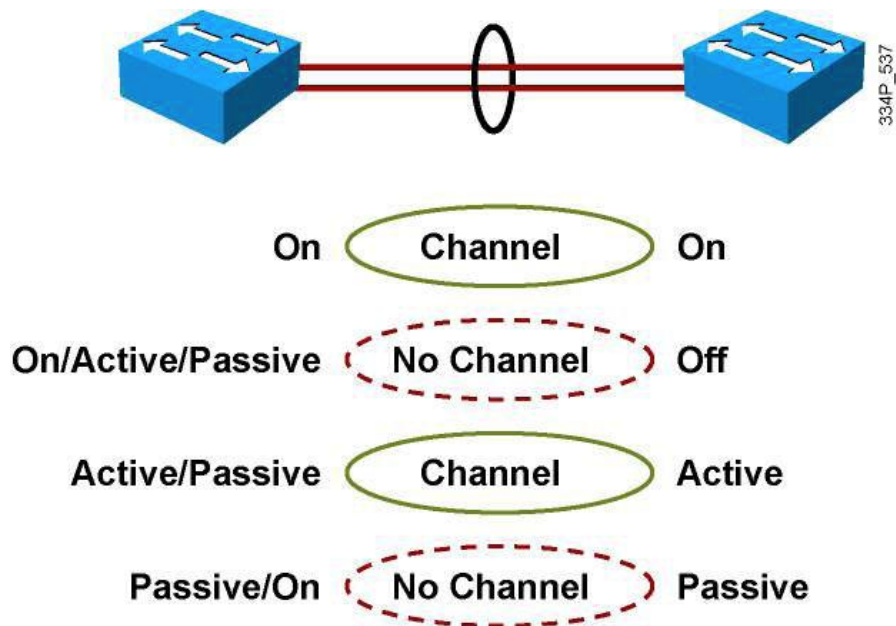
- **Port Aggregation Protocol (PAgP)** is a Cisco-proprietary protocol that aids in the automatic creation of Fast EtherChannel links.
 - When an EtherChannel link is configured using PAgP, PAgP packets are sent between Fast EtherChannel-capable ports to negotiate the forming of a channel.
 - When PAgP identifies matched Ethernet links, it groups the links into an EtherChannel. Spanning tree adds the EtherChannel as a single bridge port.
- **Link Aggregation Control Protocol (LACP)** is part of an IEEE specification (802.3ad) that also enables several physical ports to be bundled together to form an EtherChannel.
 - LACP enables a switch to negotiate an automatic bundle by sending LACP packets to the peer.
 - It performs a similar function as PAgP with Cisco EtherChannel.
 - Because LACP is an IEEE standard, you can use it to facilitate EtherChannels in mixed-switch environments. In a Cisco environment, both protocols are supported.

PAgP Modes



Mode	Purpose
Auto	Places an interface in a passive negotiating state in which the interface responds to the PAgP packets that it receives but does not initiate PAgP negotiation (default).
Desirable	Places an interface in an active negotiating state in which the interface initiates negotiations with other interfaces by sending PAgP packets. Interfaces configured in the “on” mode do not exchange PAgP packets.
On	Forces the interface to channel without PAgP.
Non-silent	If a switch is connected to a partner that is PAgP-capable, configure the switch interface for non-silent operation. The non-silent keyword is always used with the auto or desirable mode. If you do not specify non-silent with the auto or desirable mode, silent is assumed. The silent setting is for connections to file servers or packet analyzers; this setting enables PAgP to operate, to attach the interface to a channel group, and to use the interface for transmission.

LACP Modes



Mode	Purpose
Passive	Places a port in a passive negotiating state. In this state, the port responds to the LACP packets that it receives but does not initiate LACP packet negotiation (default).
Active	Places a port in an active negotiating state. In this state, the port initiates negotiations with other ports by sending LACP packets.
On	Forces the interface to the channel without PAgP or LACP.

Configuring EtherChannel

- **Step 1.** Specify the interfaces that will compose the EtherChannel group. Using the range commands enables you to select several interfaces and configure them all together. A good practice is to start by shutting down these interfaces, so that incomplete configuration will not start to create activity on the link:

```
Switch(config)# interface range interface_type [interface_range]
```

- **Step 2.** Specify the channeling protocol to be used. This command is not applicable to all Catalyst platforms. You can also specify the channeling protocol at Step 3:

```
Switch(config-if-range)# channel-protocol {pagp | lacp}
```

- **Step 3.** Create the port-channel interface, if necessary, and assign the specified interfaces to it:

```
Switch(config-if-range)# channel-group number mode {active | on | {auto
    [non-silent]} | {desirable [non-silent]} | passive
```


- **Step 4.** Specify the port-channel interface. When in the interface configuration mode, you can configure additional parameters. The physical interfaces will inherit these parameters. When this configuration is complete, you can reenabling the physical ports in the EtherChannel bundle:

```
Switch(config)# interface port-channel number
```

```
Switch(config-if)# interface parameters
```


Example: EtherChannel Configuration

```
Switch(config)# interface fastethernet 0/23
Switch(config-if)# channel-group 2 mode active
Switch(config)# interface fastethernet 0/24
Switch(config-if)# channel-group 2 mode active
Switch(config)# interface port-channel 2
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport trunk native VLAN 99
Switch(config-if)# switchport trunk allowed VLAN 2,3,99
```



Remote Switch configuration

```
RSwitch(config)# interface fastethernet 0/23
RSwitch(config-if)# channel-group 5 mode on
RSwitch(config)# interface fastethernet 0/24
RSwitch(config-if)# channel-group 5 mode on
RSwitch(config)# interface port-channel 5
RSwitch(config-if)# switchport mode trunk
RSwitch(config-if)# switchport trunk native VLAN 99
```

Verifying EtherChannel (1)

- You can use several commands to verify an EtherChannel configuration. On any physical interface member of an EtherChannel bundle, the **show interfaces *interface_id* etherchannel** command provides information on the role of the interface in the EtherChannel.
- Interface FastEthernet 0/24 below is part of EtherChannel bundle 1.
- The protocol for this EtherChannel is LACP.

```
Switch# show interfaces fa0/24 etherchannel
Port state = Up Sngl-port-Bndl Mstr Not-in-Bndl
Channel group = 1          Mode = Active    Gchange = -
Port-channel = null       GC = -          Pseudo port-channel = Po1
Port index = 0            Load = 0x00    Protocol = LACP
```

Verifying EtherChannel (2)

- The **show etherchannel *number* port-channel** command can be used to display information about a specific port-channel.
- Below Port-channel 1 consists of two physical ports, Fa0/23 and Fa0/24.
- It uses LACP in active mode.
- It is properly connected to another switch with a compatible configuration. This is why the port-channel is said to be in use.

```
Switch# show etherchannel 1 port-channel
Port-channels in the group:
-----
Port-channel: Po7          (Primary Aggregator)
Age of the Port-channel = 195d:03h:10m:44s
Logical slot/port = 0/1    Number of ports = 2
Port state = Port-channel Ag-Inuse
Protocol = LACP

Ports in the Port-channel:
Index   Load   Port      EC state   No of bits
-----+-----+-----+-----+-----
0       55     fa0/23    Active     4
1       45     fa0/24    Active     4
```

Verifying EtherChannel (3)

- When several port-channel interfaces are configured on the same device, the **show etherchannel summary** command is useful for displaying one-line information per port-channel.
- As shown below; the switch has three EtherChannels configured: Groups 2 and 7 use LACP and Group 9 uses PAgP. Each EtherChannel has the member interfaces listed. All three groups are Layer 2 EtherChannels and are all in use (**SU** next to the port-channel number).

```
Switch# 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
Number of channel-groups in use: 2
Number of aggregators: 2
Group      Port-channel  Protocol    Ports
-----+-----+-----+-----
2          Po2 (SU)      LACP        g0/49 (P) g0/50 (P) g0/51 (P) g0/52 (P)
7          Po7 (SU)      LACP        g0/47 (P) g0/48 (P)
9          Po9 (SU)      PAgP        g0/8 (P)  g0/9 (P)
```

Verifying EtherChannel (4)

- The **show running-config interface *interface_id*** command displays sections of your configuration relevant to EtherChannel. The interface argument can be physical or logical.

```
Switch# show running-config interface g0/48
```

```
Building configuration...
```

```
Current configuration : 154 bytes
```

```
interface GigabitEthernet0/48
```

```
switchport access vlan 41
```

```
switchport trunk encapsulation dot1q
```

```
switchport mode trunk
```

```
channel-group 7 mode active
```

```
Switch# show running-config interface port-channel 7
```

```
Building configuration...
```


```
Current configuration : 92 bytes
```

```
interface Port-channel7
```

```
switchport trunk encapsulation dot1q
```

```
switchport mode trunk
```

EtherChannel Load Balancing



```
switch(config)# port-channel load-balance type
```

```
switch# show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
src-dst-ip
```

EtherChannel Load Balancing Example

- Here the EtherChannel load-balancing mechanism is configured to use source and destination IP address pairs.
- This rule is applied to IPv4 and IPv6 traffic, whereas the non-IP load-balancing mechanism uses source and destination MAC address pairs.
- It was observed that with source-destination IP load balancing, the balancing ends up more like 70-30 on the links!

```
Switch(config)# port-channel load-balance src-dst-ip
Switch(config)# exit
Switch# show etherchannel load-balance
EtherChannel Load-Balancing Configuration:
src-dst-ip
EtherChannel Load-Balancing Addresses Used Per-Protocol:
Non-IP: Source XOR Destination MAC address
IPv4: Source XOR Destination IP address
IPv6: Source XOR Destination IP address
```

Chapter 1 Summary

- A VLAN is a logical grouping of switch ports independent of physical location. Local VLANs are now recommended over end-to-end VLAN implementations.
- A trunk is a Layer 2 point-to-point link between networking devices carry the traffic of multiple VLANs.
- ISL and 802.1Q are the two trunking protocols that can connect two switches.
- VTP is used to distribute and synchronize information about VLANs configured throughout a switched network.
- VTP pruning helps to stop flooding of unnecessary traffic on trunk links.
- Device communication within the same VLAN can be fine-tuned using pVLANs. A pVLAN is associated to a primary VLAN, and then mapped to one or several ports. A primary VLAN can map to one isolated and several community VLANs. pVLANs can span across several switches using regular 802.1q trunks or pVLAN trunks.
- Use EtherChannel by aggregating individual, similar links between switches. EtherChannel can be dynamically configured between switches using either the Cisco-proprietary PAgP or the IEEE 802.3ad LACP. EtherChannel load balances traffic over all the links in the bundle. The method that is chosen directly impacts the efficiency of this load-balancing mechanism.

Chapter 1 Labs

- **SW-LAB-1**
 - VTP, VLANs, Trunking
 - Etherchannel

Q&A