Cisco

Exam 350-001

CCIE Routing and Switching Written Exam v4.0

Version: 15.0

[Total Questions: 572]

Topic 1, Implement Layer 2 Technologies

Question No : 1 - (Topic 1)

Which statement is true about loop guard?

A. Loop guard only operates on interfaces that are considered point-to-point by the spanning tree.
B. Loop guard only operates on root ports.
C. Loop guard only operates on designated ports.
D. Loop guard only operates on edge ports.

Answer: A

Explanation:

Understanding How Loop Guard Works

Unidirectional link failures may cause a root port or alternate port to become designated as root if BPDUs are absent. Some software failures may introduce temporary loops in the network. Loop guard checks if a root port or an alternate root port receives BPDUs. If the port is receiving BPDUs, loop guard puts the port into an inconsistent state until it starts receiving BPDUs again. Loop guard isolates the failure and lets spanning tree converge to a stable topology without the failed link or bridge.

You can enable loop guard per port with the set span tree guard loop command.

Note When you are in MST mode, you can set all the ports on a switch with the set span tree global-defaults loop-guard command.

When you enable loop guard, it is automatically applied to all of the active instances or VLANs to which that port belongs. When you disable loop guard, it is disabled for the specified ports. Disabling loop guard moves all loop-inconsistent ports to the listening state.

If you enable loop guard on a channel and the first link becomes unidirectional, loop guard blocks the entire channel until the affected port is removed from the channel. Figure 8-6 shows loop guard in a triangle switch configuration.

Figure 8-6 Triangle Switch Configuration with Loop Guard

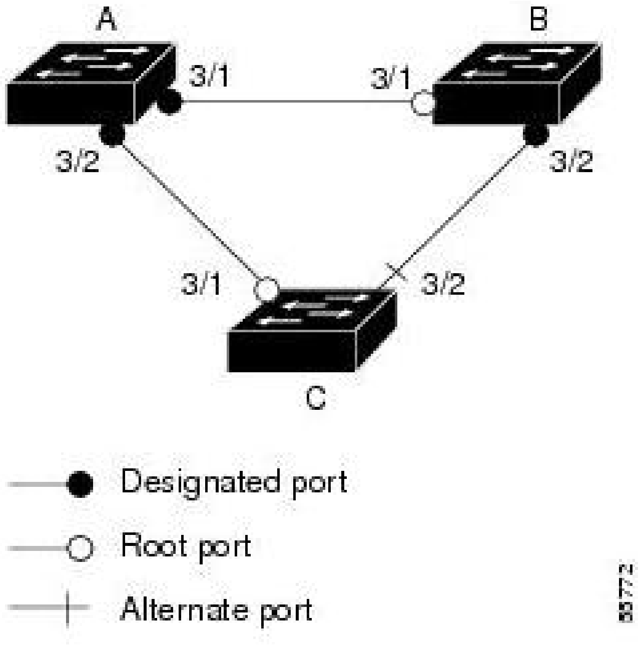

Figure 8-6 illustrates the following configuration:

Switches A and B are distribution switches.

Switch C is an access switch.

Loop guard is enabled on ports 3/1 and 3/2 on Switches A, B, and C.

Use loop guard only in topologies where there are blocked ports. Topologies that have no blocked ports, which are loop free, do not need to enable this feature. Enabling loop guard on a root switch has no effect but provides protection when a root switch becomes a nonroot switch.

Follow these guidelines when using loop guard:

Do not enable loop guard on PortFast-enabled or dynamic VLAN ports.

Do not enable PortFast on loop guard-enabled ports.

Do not enable loop guard if root guard is enabled.

Do not enable loop guard on ports that are connected to a shared link.

Note: We recommend that you enable loop guard on root ports and alternate root ports on access switches.

Loop guard interacts with other features as follows:

Loop guard does not affect the functionality of UplinkFast or BackboneFast.

Root guard forces a port to always be designated as the root port. Loop guard is effective only if the port is a root port or an alternate port. Do not enable loop guard and root guard on a port at the same time.

PortFast transitions a port into a forwarding state immediately when a link is established. Because a PortFast-enabled port will not be a root port or alternate port, loop guard and

PortFast cannot be configured on the same port. Assigning dynamic VLAN membership for the port requires that the port is PortFast enabled. Do not configure a loop guard-enabled port with dynamic VLAN membership.

If your network has a type-inconsistent port or a PVID-inconsistent port, all BPDUs are dropped until the misconfiguration is corrected. The port transitions out of the inconsistent state after the message age expires. Loop guard ignores the message age expiration on type-inconsistent ports and PVID-inconsistent ports. If the port is already blocked by loop guard, misconfigured BPDUs that are received on the port make loop guard recover, but the port is moved into the type-inconsistent state or PVID-inconsistent state.

In high-availability switch configurations, if a port is put into the blocked state by loop guard, it remains blocked even after a switchover to the redundant supervisor engine. The newly activated supervisor engine recovers the port only after receiving a BPDU on that port.

Loop guard uses the ports known to spanning tree. Loop guard can take advantage of logical ports provided by the Port Aggregation Protocol (PAgP). However, to form a channel, all the physical ports grouped in the channel must have compatible configurations. PAgP enforces uniform configurations of root guard or loop guard on all the physical ports to form a channel.

These caveats apply to loop guard:

–Spanning tree always chooses the first operational port in the channel to send the BPDUs. If that link becomes unidirectional, loop guard blocks the channel, even if other links in the channel are functioning properly.

–If a set of ports that are already blocked by loop guard are grouped together to form a channel, spanning tree loses all the state information for those ports and the new channel port may obtain the forwarding state with a designated role.

–If a channel is blocked by loop guard and the channel breaks, spanning tree loses all the state information.

The individual physical ports may obtain the forwarding state with the designated role, even if one or more of the links that formed the channel are unidirectional.

You can enable UniDirectional Link Detection (UDLD) to help isolate the link failure. A loop may occur until

UDLD detects the failure, but loop guard will not be able to detect it.

Loop guard has no effect on a disabled spanning tree instance or a VLAN.

Question No : 2 - (Topic 1)

Which command is used to enable EtherChannel hashing for Layer 3 IP and Layer 4 port-based CEF?

A. mpls ip cef
B. port-channel ip cef
C. mpls ip port-channel cef
D. port-channel load balance
E. mpls ip load-balance
F. ip cef EtherChannel channel-id XOR L4
G. ip cef connection exchange

Answer: D

Explanation:

Port-channel load balance is normally used for enable etherchannel hashing for Layer 3 IP and Layer 4 port based CEF.

Question No : 3 - (Topic 1)

Which two options are contained in a VTP subset advertisement? (Choose two.)

A. Followers field
B. MD5 digest
C. VLAN information
D. Sequence number

Answer: C,D

Explanation:

Subset Advertisements

When you add, delete, or change a VLAN in a Catalyst, the server Catalyst 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 in order to advertise all the VLANs.

Subset Advertisement Packet Format

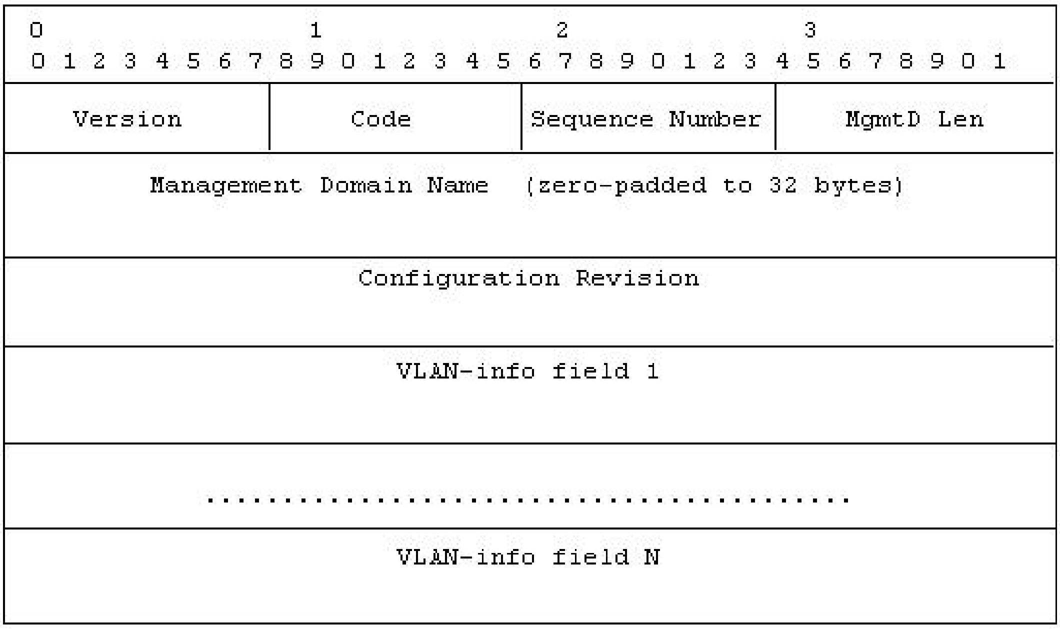

0 1 2 3 4 5 6 7	8 9 0 1 2 3 4 5	6 7 8 9 0 1 2 3	4 5 6 7 8 9 0 1
Version	Code	Sequence Number	MgmtD Len
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision			
VLAN-info field 1			
....			
VLAN-info field N			

Bit ruler above the first column: 0 (bits 0–9), 1 (bits 10–19), 2 (bits 20–29), 3 (bits 30–31).

This formatted example shows that each VLAN information field contains information for a different VLAN. It is ordered so that lowered-valued ISL VLAN IDs occur first:

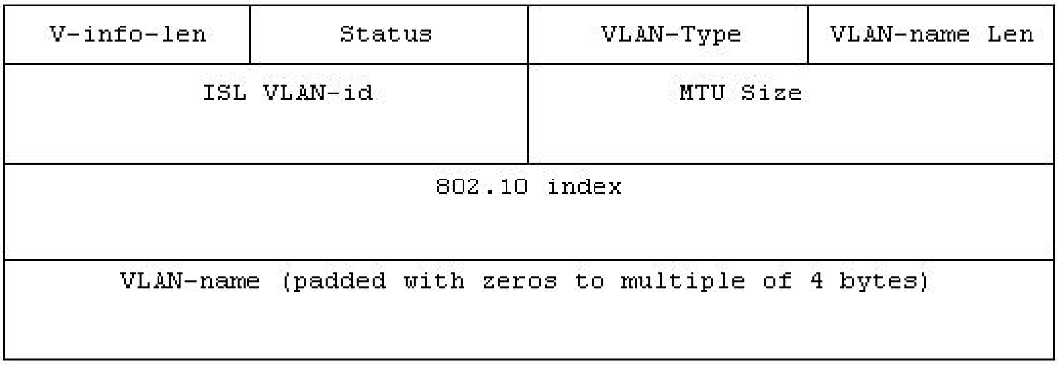

V-info-len	Status	VLAN-Type	VLAN-name Len
ISL VLAN-id		MTU Size	
802.10 index			
VLAN-name (padded with zeros to multiple of 4 bytes)			

Most of the fields in this packet are easy to understand. These are two clarifications:
Code — The format for this is 0x02 for subset advertisement.
Sequence number — This is the sequence of the packet in the stream of packets that follow a summary advertisement. The sequence starts with 1.

Advertisement Requests
A switch needs 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. This is an example:

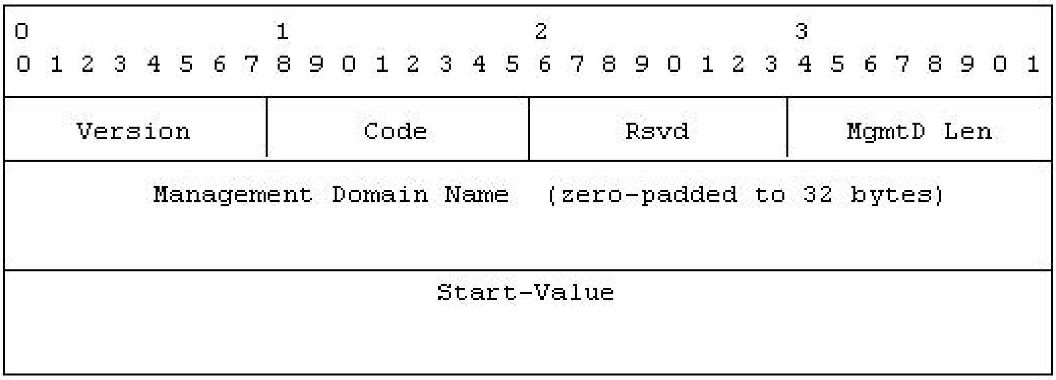

0	1	2	3
0 1 2 3 4 5 6 7	8 9 0 1 2 3 4 5	6 7 8 9 0 1 2 3	4 5 6 7 8 9 0 1
Version	Code	Rsvd	MgmtD Len
Management Domain Name (zero-padded to 32 bytes)			
Start-Value			

Code—The format for this is 0x03 for an advertisement request.

Start-Value—This is used in cases in which there are several subset advertisements. If the first (n) subset advertisement has been received and the subsequent one (n+1) has not been received, the Catalyst only requests advertisements from the (n+1)th one.

Reference

http://www.cisco.com/en/US/tech/tk389/tk689/technologies_tech_note09186a0080094c52.shtml

Question No : 4 - (Topic 1)

Which two statements are true about traffic shaping? (Choose two.)

A. Out-of-profile packets are queued.
B. It causes TCP retransmits.
C. Marking/remarking is not supported.

D. It does not respond to BECN and ForeSight Messages.
E. It uses a single/two-bucket mechanism for metering.

Answer: A,C

Question No : 5 - (Topic 1)

Which three options are features of VTP version 3? (Choose three.)

A. VTPv3 supports 8K VLANs.
B. VTPv3 supports private VLAN mapping.
C. VTPv3 allows for domain discovery.
D. VTPv3 uses a primary server concept to avoid configuration revision issues.
E. VTPv3 is not compatible with VTPv1 or VTPv2.
F. VTPv3 has a hidden password option.

Answer: B,D,F

Explanation:

Key Benefits of VTP Version 3

Much work has gone into improving the usability of VTP version 3 in three major areas:

The new version of VTP offers better administrative control over which device is allowed to update other devices' view of the VLAN topology. The chance of unintended and disruptive changes is significantly reduced, and availability is increased. The reduced risk of unintended changes will ease the change process and help speed deployment.

Functionality for the VLAN environment has been significantly expanded. Two enhancements are most beneficial for today's networks:

– In addition to supporting the earlier ISL VLAN range from 1 to 1001, the new version supports the whole IEEE 802.1Q VLAN range up to 4095.

– In addition to supporting the concept of normal VLANs, VTP version 3 can transfer information regarding Private VLAN (PVLAN) structures.

The third area of major improvement is support for databases other than VLAN (for example, MST).

Brief Background on VTP Version 1 and VTP Version 2

VTP version 1 was developed when only 1k VLANs where available for configuration. A tight internal coupling of the VLAN implementation, the VLAN pruning feature, and the VTP function itself offered an efficient means of implementation. It has proved in the field to reliably support Ethernet, Token Ring, and FDDI networks via VTP.

The use of consistent VLAN naming was a requirement for successful use of VMPS (Vlan Membership Policy Server). VTP ensures the consistency of VLAN names across the VTP domain. Most VMPS implementations are likely to be migrated to a newer, more flexible and feature-rich method.

To add support for Token Ring, VTP version 1 was enhanced and called VTP version 2. Certain other minor changes and enhancements were also added at this time.

The functional base in VTP version 3 is left unchanged from VTP version 2, so backward compatibility is built in. It is possible, on a per link basis, to automatically discover and support VTP version 2 devices.

VTP version 3 adds a number of enhancements to VTP version 1 and VTP version 2:
Support for a structured and secure VLAN environment (Private VLAN, or PVLAN)
Support for up to 4k VLANs
Feature enhancement beyond support for a single database or VTP instance
Protection from unintended database overrides during insertion of new switches
Option of clear text or hidden password protection
Configuration option on a per port base instead of only a global scheme
Optimized resource handling and more efficient transfer of information
These new requirements made a new code foundation necessary. The design goal was to make VTP version 3 a versatile vehicle. This was not only for the task of transferring a VLAN DB but also for transferring other databases-for example, the MST database.

Reference
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/solution_guide_c78_508010.html

Question No : 6 - (Topic 1)

Which three options are considered in the spanning-tree decision process? (Choose three.)

A. Lowest root bridge ID
B. Lowest path cost to root bridge
C. Lowest sender bridge ID
D. Highest port ID
E. Highest root bridge ID
F. Highest path cost to root bridge

Answer: A,B,C

Explanation:

Configuration bridge protocol data units (BPDUs) are sent between switches for each port. Switches use s four step process to save a copy of the best BPDU seen on every port. When a port receives a better BPDU, it stops sending them. If the BPDUs stop arriving for 20 seconds (default), it begins sending them again.

Step 1 Lowest Root Bridge ID (BID)
Step 2 Lowest Path Cost to Root Bridge
Step 3 Lowest Sender BID
Step 4 Lowest Port ID

Reference
Cisco General Networking Theory Quick Reference Sheets

Question No : 7 - (Topic 1)

In 802.1s, how is the VLAN to instance mapping represented in the BPDU?

A. The VLAN to instance mapping is a normal 16-byte field in the MST BPDU.
B. The VLAN to instance mapping is a normal 12-byte field in the MST BPDU.
C. The VLAN to instance mapping is a 16-byte MD5 signature field in the MST BPDU.
D. The VLAN to instance mapping is a 12-byte MD5 signature field in the MST BPDU.

Answer: C

Explanation:

MST Configuration and MST Region
Each switch running MST in the network has a single MST configuration that consists of these three attributes:
1. An alphanumeric configuration name (32 bytes)
2. A configuration revision number (two bytes)
3. A 4096-element table that associates each of the potential 4096 VLANs supported on the chassis to a given instance.

In order to be part of a common MST region, a group of switches must share the same