HP0-Y43 HP ExpertONE Certification Exam

Number: EnsurePass
Passing Score: 740
Time Limit: 120 min
File Version: 13.01

Vendor: HP
Exam Code: HP0-Y43
Exam Name: Implementing HP Network Infrastructure Solutions Exam
Version: 13.01

Contact us:
If you have any suggestions or any questions about our product,please feel free to contact us: support@ensurepass.com

About Products:
Free update is available within 180 days after your purchase.
Please login your user center and download the latest product anytime.
PS:Ensure you can pass the exam,please check the latest product in 2-3 days before the exam again.

Sections

1. IMC
2. IPv6
3. High Availability
4. MSTP
5. Multicast
6. OSPF
7. Other
8. QOS
9. Security
10. VLANs

Exam A

QUESTION 1
A customer requires an HP FlexCampus solution with a core that scales to 40/100G. Which HP switch fabric meets this need?

A. the 7500's CLOS switch fabric
B. the 10500's CLOS switch fabric
C. the 7500's crossbar switch fabric
D. the 10500's crossbar switch fabric

Correct Answer: B
Section: Other
Explanation

Explanation/Reference:
HP 7500 Switch Series
Key features
Versatile, high-performance modular switches
Enterprise LAN core, aggregation, and edge
Extensive switching and routing, IPv6, MPLS
Advanced functionality with service modules
Robust network and service virtualization

HP 10500 Switch Series
Key features
Advanced, next-generation **CLOS architecture**
More than 11 terabits-per-second switching capacity
Feature-rich, with IPv6 and MPLS functionality
HP IRF technology virtualizes up to four chassis
Ultra-high 1/10/40 GbE density; **100 GbE ready**

QUESTION 2
What is the role of neighbor solicitation (NS) messages in the autoconfiguration of an IPv6 address?

A. An IPv6 node sends an NS message to inform a node undergoing autoconfiguration that it is already using a particular address.
B. An IPv6 node sends an NS message for its tentative address to determine whether another node is using it.
C. An IPv6 node sends an NS message for the global prefix to prompt other IPv6 nodes to advertise the addresses that they are using on that prefix.
D. An IPv6 node sends an NS message to prompt an IPv6 router on the link to advertise the global prefixes associated with the link immediately.

Correct Answer: B
Section: IPv6

Explanation

Explanation/Reference:
Duplicate address detection

The assignment of a unicast IPv6 address to an interface involves an internal test for the uniqueness of that address using Neighbor Solicitation and Neighbor Advertisement (ICMPv6 type 135 and 136) messages. While in the process of establishing uniqueness an address has a tentative state.
The node joins the solicited-node multicast address for the tentative address (if not already done so) and sends neighbor solicitations, with the tentative address as target address and the unspecified address (::/128) as source address. The node also joins the all-hosts multicast address ff02::1, so it will be able to receive Neighbor Advertisements.
If a node receives a neighbor solicitation with its own tentative address as the target address, then that address is not unique. The same is true if the node receives a neighbor advertisement with the tentative address as the source of the advertisement. Only after having successfully established that an address is unique may it be assigned and used by an interface.

QUESTION 3
Which switch is best suited to act at the edge of a medium to large HP FlexFabric solution?

A. 10500
B. 5500
C. 9500
D. 5830

Correct Answer: D
Section: Other
Explanation

Explanation/Reference:
Technical white paper
HP FlexFabric Reference Architecture Overview

Customers looking to reduce cost and complexity can implement a two-tier collapsed network design that completely eliminates a dedicated aggregation layer. These designs leverage HP Virtual Connect or HP 58x0/59x0 series switches at the server edge along with highly scalable HP12500 series core switches as a collapsed core/aggregation layer. These flat network designs help ensure direct-flight server-to-server performance while dramatically reducing network port counts. A two-tier collapsed design also simplifies and streamlines network management,and reduces capital expense and energy consumption.

http://www.hp.com/hpinfo/newsroom/press_kits/2012/convergedcloud2012/HPN_FlexFabric_Whitepaper.pdf

QUESTION 4
How can a high density of ports and high throughput at the core of an HP FlexNetwork save customers money?

A. Intelligence is offloaded from the edge switches, enabling customers to save money on the most numerous switches in their solutions.

B. The customer can combine the data center and campus LAN core into a single entity, reducing power and cooling costs.
C. The customer no longer needs to deploy modular switches at the distribution level and the edge, deploying more cost-effective stackable switches instead.
D. The architecture can be simplified, eliminating expensive distribution devices and reducing power and cooling costs.

Correct Answer: D
Section: Other
Explanation

Explanation/Reference:

QUESTION 5
View the exhibit. The exhibit shows a network with HP 5400 zl and 8200 zl switches throughout the core and edge. What is one advantage of implementing routing on edge switches?

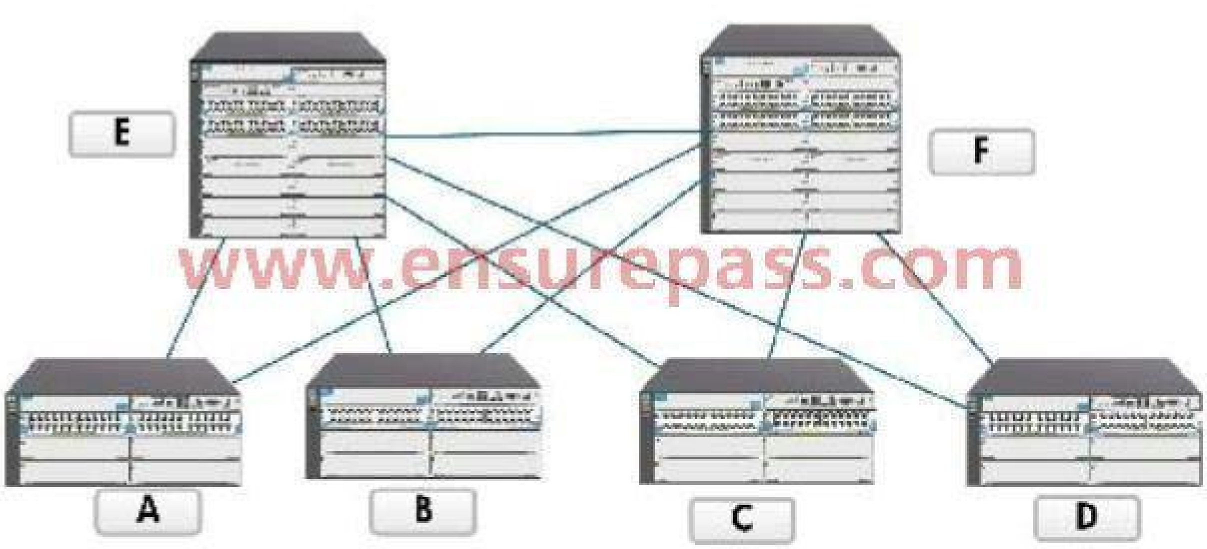

A. Typically, it is easier to implement user-based VLAN assignments.
B. The topology has higher redundancy because edge switches can take over routing roles when necessary.
C. Typically, the network can use fewer total VLANs when edge switches implement routing.
D. Typically, it is easier to ensure that routed links between edge and core switches are fully utilized.

Correct Answer: D
Section: Other
Explanation

Explanation/Reference:
Remember

QUESTION 6
A company has a network that includes HP 5800 and 12500 switches. Usage demands on the company's FTP servers have increased, causing performance issues during peak usage times. While analyzing link utilization, the network administrator noticed that the most heavily utilized links experience bursts of congestion, causing them to drop traffic. The links then experience brief periods of low utilization followed by another burst of congestion. This pattern continues periodically throughout the peak utilization time.

What should the network administrator do to attempt to create a more efficient traffic pattern on these links?

A. Configure an outbound traffic policing policy on the ports in question, setting the CIR at about sixty percent of the ports' capacity and the PIR at abou eighty percent.
B. Apply Weighted Fair Queuing (WFQ) or Weighted Round Robin (WRR) scheduling in preference to Strict Priority (SP) scheduling on the ports in question.
C. Configure inbound traffic policing policies on ports at the core, setting the CIR at about sixty percent of the ports' capacity and the PIR at about eighty percent. Apply outbound generic traffic shaping (GTS) on ports facing the core ports, setting the CIR equal to the CIR on the core ports.
D. Apply a WRED table to the ports in question, optionally adjusting the table values to drop lower priority traffic first.

Correct Answer: D
Section: QOS
Explanation

Explanation/Reference:
WRED proceeds in this order when a packet arrives:

- Calculation of the average queue size.
- The arriving packet is queued only if the average queue size is below the minimum queue threshold.
- Depending on the packet drop probability the packet is either dropped or queued if the average queue size is between the minimum and maximum queue threshold.
- The packet is automatically dropped if the average queue size is greater than the maximum threshold.

QUESTION 7
View the exhibits.

Exhibit 1
The frame has these characteristics:

VLANID = 3

802.1p = 3
DSCP=32
Source IP address = 10.1.1.5

Exhibit 2

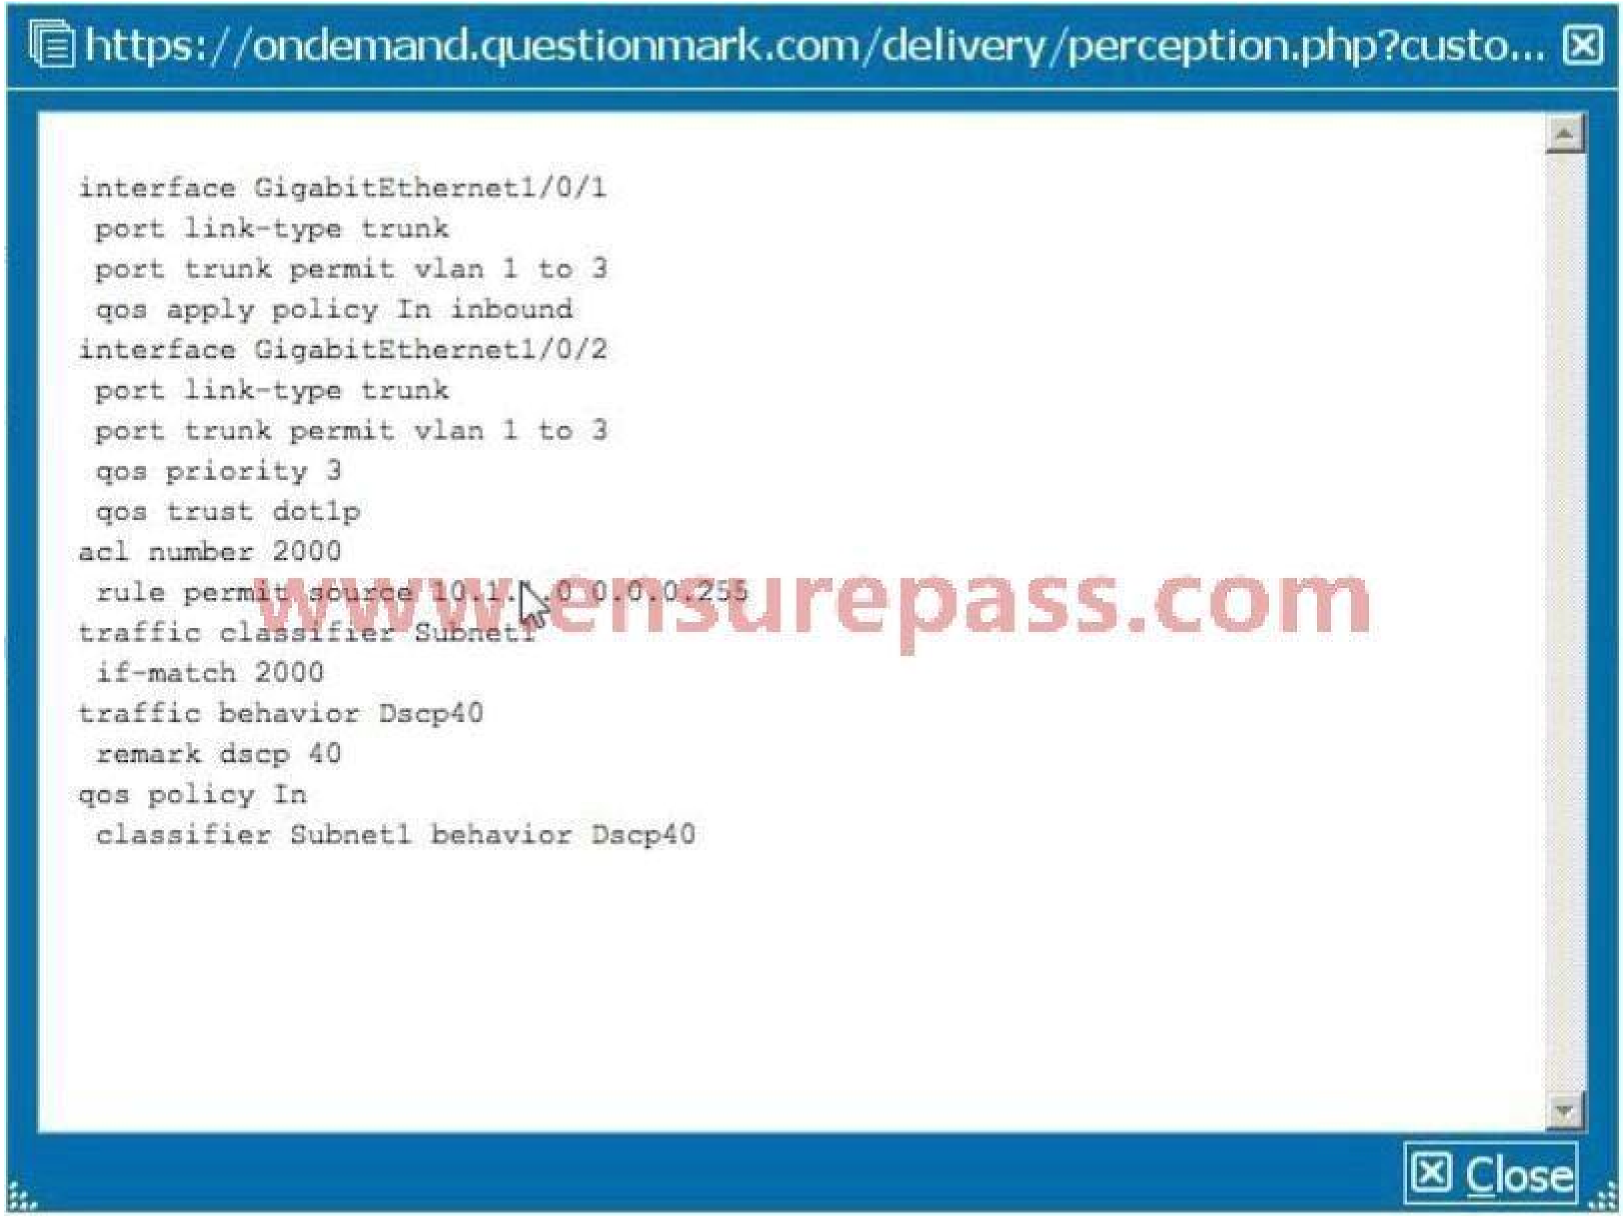

https://ondemand.questionmark.com/delivery/perception.php?custo...
☒ Close

```
interface GigabitEthernet1/0/1
 port link-type trunk
 port trunk permit vlan 1 to 3
 qos apply policy In inbound
interface GigabitEthernet1/0/2
 port link-type trunk
 port trunk permit vlan 1 to 3
 qos priority 3
 qos trust dot1p
acl number 2000
 rule permit source 10.1.[illegible].0 0.0.0.255
traffic classifier Subnet1
 if-match 2000
traffic behavior Dscp40
 remark dscp 40
qos policy In
 classifier Subnet1 behavior Dscp40
```

The frame shown in Exhibit 1 arrives on an HP 5800 switch's Gigabit Ethernet port 1/0/1. QoS maps are at their default settings. Based on the configuration shown in Exhibit 2, to which queue is the outbound packet assigned?

A. 2
B. 3
C. 4
D. 5

Correct Answer: D
Section: QOS
Explanation

Explanation/Reference:
In accordance with the table below: DSCP == 40 => TOS == 5

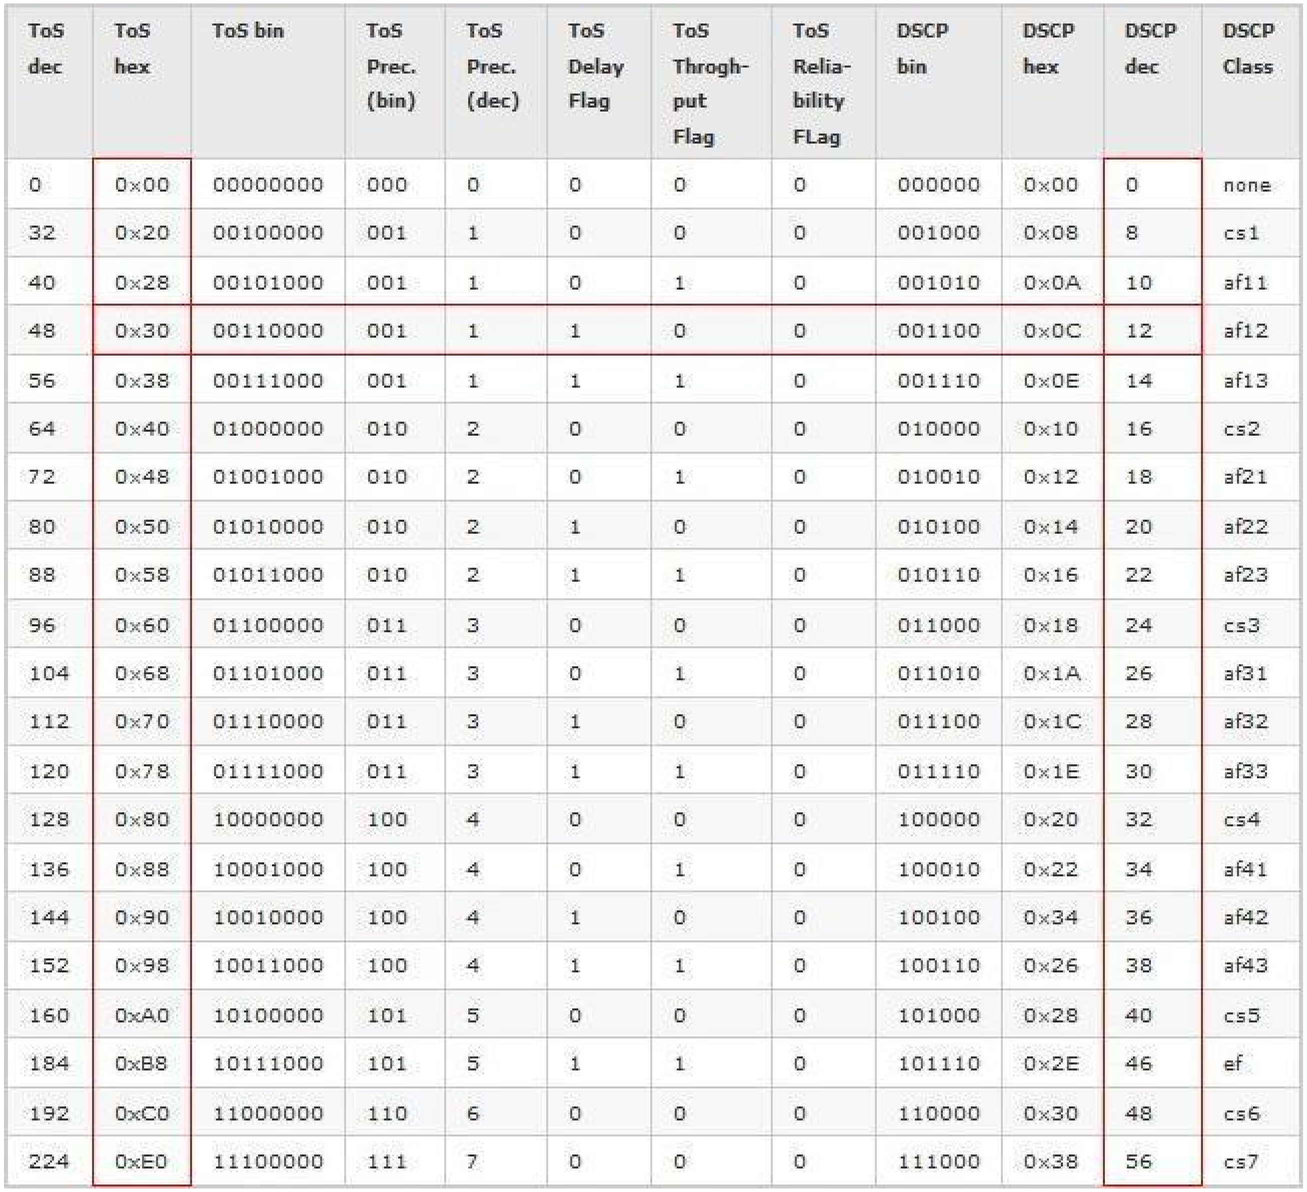

ToS dec	ToS hex	ToS bin	ToS Prec. (bin)	ToS Prec. (dec)	ToS Delay Flag	ToS Throgh-put Flag	ToS Reliability FLag	DSCP bin	DSCP hex	DSCP dec	DSCP Class
0	0x00	00000000	000	0	0	0	0	000000	0x00	0	none
32	0x20	00100000	001	1	0	0	0	001000	0x08	8	cs1
40	0x28	00101000	001	1	0	1	0	001010	0x0A	10	af11
48	0x30	00110000	001	1	1	0	0	001100	0x0C	12	af12
56	0x38	00111000	001	1	1	1	0	001110	0x0E	14	af13
64	0x40	01000000	010	2	0	0	0	010000	0x10	16	cs2
72	0x48	01001000	010	2	0	1	0	010010	0x12	18	af21
80	0x50	01010000	010	2	1	0	0	010100	0x14	20	af22
88	0x58	01011000	010	2	1	1	0	010110	0x16	22	af23
96	0x60	01100000	011	3	0	0	0	011000	0x18	24	cs3
104	0x68	01101000	011	3	0	1	0	011010	0x1A	26	af31
112	0x70	01110000	011	3	1	0	0	011100	0x1C	28	af32
120	0x78	01111000	011	3	1	1	0	011110	0x1E	30	af33
128	0x80	10000000	100	4	0	0	0	100000	0x20	32	cs4
136	0x88	10001000	100	4	0	1	0	100010	0x22	34	af41
144	0x90	10010000	100	4	1	0	0	100100	0x34	36	af42
152	0x98	10011000	100	4	1	1	0	100110	0x26	38	af43
160	0xA0	10100000	101	5	0	0	0	101000	0x28	40	cs5
184	0xB8	10111000	101	5	1	1	0	101110	0x2E	46	ef
192	0xC0	11000000	110	6	0	0	0	110000	0x30	48	cs6
224	0xE0	11100000	111	7	0	0	0	111000	0x38	56	cs7

Priority trust mode on a port

The priority trust mode on a port determines which priority is used for priority mapping table lookup. priority was introduced so that you can use it for priority mapping in addition to the priority fields car in packets. The HP 5800 Switch Series and 5820X Switch Series provide the following priority trust mo

- **dot1p**—Uses the 802.1p priority carried in packets for priority mapping.

Table 3 Priority mapping results of trusting the 802.1p priority (when the default dot1p-lp priority mapping table is used)

802.1p priority carried in packets	Local precedence	Queue ID
0	2	2
1	0	0
2	1	1
3	3	3
4	4	4
5	5	5
6	6	6
7	7	7

NOTE:

When the 802.1p priority carried in packets is trusted, the port priority is used for priority mapping for packets which do not carry VLAN tags (namely, do not carry 802.1p priorities.) The priority mapping

QUESTION 8
View the exhibit. A network administrator wants to configure an HP 5800 switch to place all incoming traffic on the Gigabit Ethernet 1/0/1 port in priority queue 3. However, one exception applies. All traffic incoming on that port that is destined to UDP port 55555 should be forwarded in priority queue 4 and marked with DSCP 32. Based on the current configurations shown in the exhibit, which further steps must the network administrator perform? (Select two.)

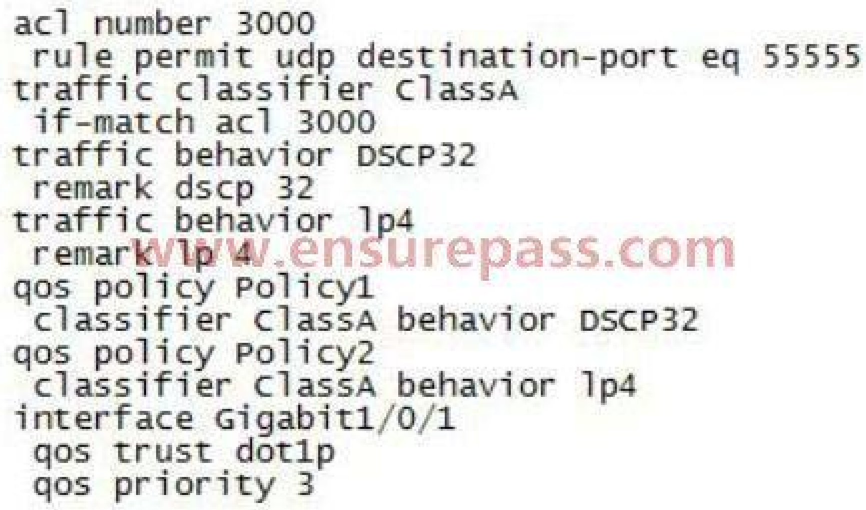

```
acl number 3000
 rule permit udp destination-port eq 55555
traffic classifier ClassA
 if-match acl 3000
traffic behavior DSCP32
 remark dscp 32
traffic behavior lp4
 remark lp 4
qos policy Policy1
 classifier ClassA behavior DSCP32
qos policy Policy2
 classifier ClassA behavior lp4
interface Gigabit1/0/1
 qos trust dot1p
 qos priority 3
```

A. Configure port Gigabit Ethernet 1/0/1 to trust DSCP
B. Apply QoS policy Policy1 as an inbound policy on port Gigabit Ethernet 1/0/1
C. Apply QoS policy Policy2 as an inbound policy on port Gigabit Ethernet 1/0/1
D. Create a OoS Ip-dscp map that maps Ip value 4 to DSCP02
E. Undo OoS trust on port Gigabit Ethernet 1/0/1

Correct Answer: BE
Section: QOS
Explanation

Explanation/Reference:

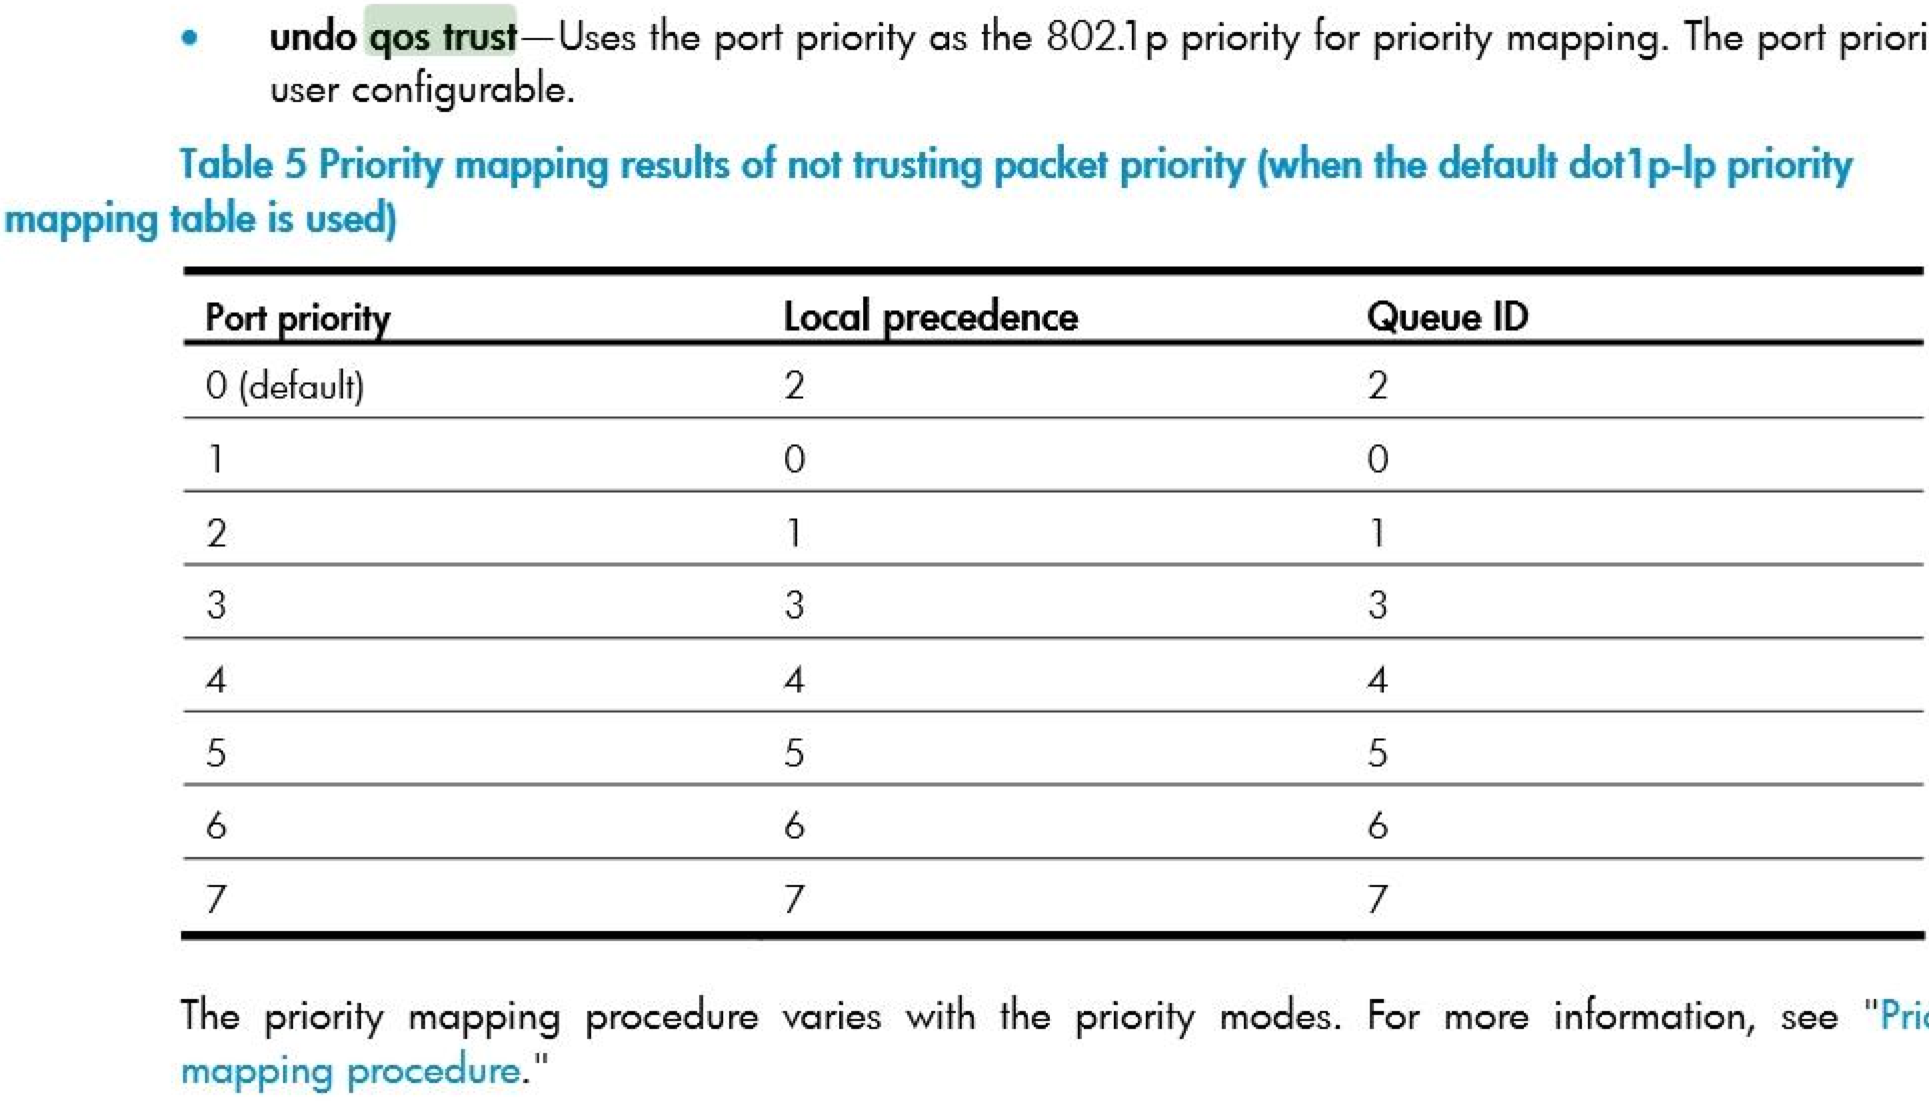

- **undo qos trust**—Uses the port priority as the 802.1p priority for priority mapping. The port priori user configurable.

Table 5 Priority mapping results of not trusting packet priority (when the default dot1p-lp priority mapping table is used)

Port priority	Local precedence	Queue ID
0 (default)	2	2
1	0	0
2	1	1
3	3	3
4	4	4
5	5	5
6	6	6
7	7	7

The priority mapping procedure varies with the priority modes. For more information, see "Pri mapping procedure."

QUESTION 9

A company has a service level agreement (SLA) with its service provider. The SLA specifies a 2 Mbps committed information rate (CIR) and 20 KB committed burst size (CBS). Lately, during peak usage times, the company has been experiencing brief periods of poor performance on its external connection (Gigabit Ethernet port 2/0/1 on an HP 5800 switch). How can the network administrator configure the Comware switch to address this problem?

A. Configure a traffic classifier that selects all traffic and a car traffic behavior that sets the CIR to 2 Mbps and the CBS to 20 KB. Create a QoS policy that maps the classifier to the action and apply this policy as an inbound policy on port Gigabit Ethernet 2/0/1.
B. Enable Strict Priority (SP) scheduling on port Gigabit Ethernet 2/0/1 and ensure that all inbound traffic is marked with the correct priority.
C. Configure generic traffic shaping (GTS) on the Gigabit Ethernet port 2/0/1, setting the CIR to 2 Mbps and the CBS to 20 KB.
D. Configure a traffic classifier that selects all traffic and a car traffic behavior that sets the CIR to 2 Mbps and the CBS to 20 KB. Create a QoS policy that maps the classifier to the action and apply this policy as an outbound policy on port Gigabit Ethernet 2/0/1.

Correct Answer: C
Section: QOS
Explanation

Explanation/Reference:
In this case you must use GTS (generic traffic shaping).
Note: Traffic shaping applies only to outbound traffic!

QUESTION 10
A network includes a mix of IGMPv2 and IGMPv3 endpoints and must support the following source specific multicasting applications:

Source: 10.1.4.2 and Group: 232.0.5.5
Source: 10.1.4.12 and Group 232.0.6.6

The network is already implementing PIM-SM and IGMPv3. In order to support these applications, the HP 10500 switches that act as routers for the endpoints in question must support another feature. Which step must the network administrator perform on each of these switches?

A. Create an SSM policy that includes 232.0.5.5 and 232.0.6.6 within its range
B. Enable IGMPv2 backward compatibility mode
C. Configure two SSM maps, each of which maps a source to its multicast group
D. Enable PIM SSM on the interfaces that connect to the endpoints

Correct Answer: C
Section: Multicast
Explanation

Explanation/Reference:
HP 10500 Switch Series IP Multicast Configuration Guide

Because of some possible restrictions, some receiver hosts on an SSM network might run IGMPv1 or IGMPv2. To provide SSM service support for these receiver hosts, configure the IGMP mapping feature on the last-hop router.
...

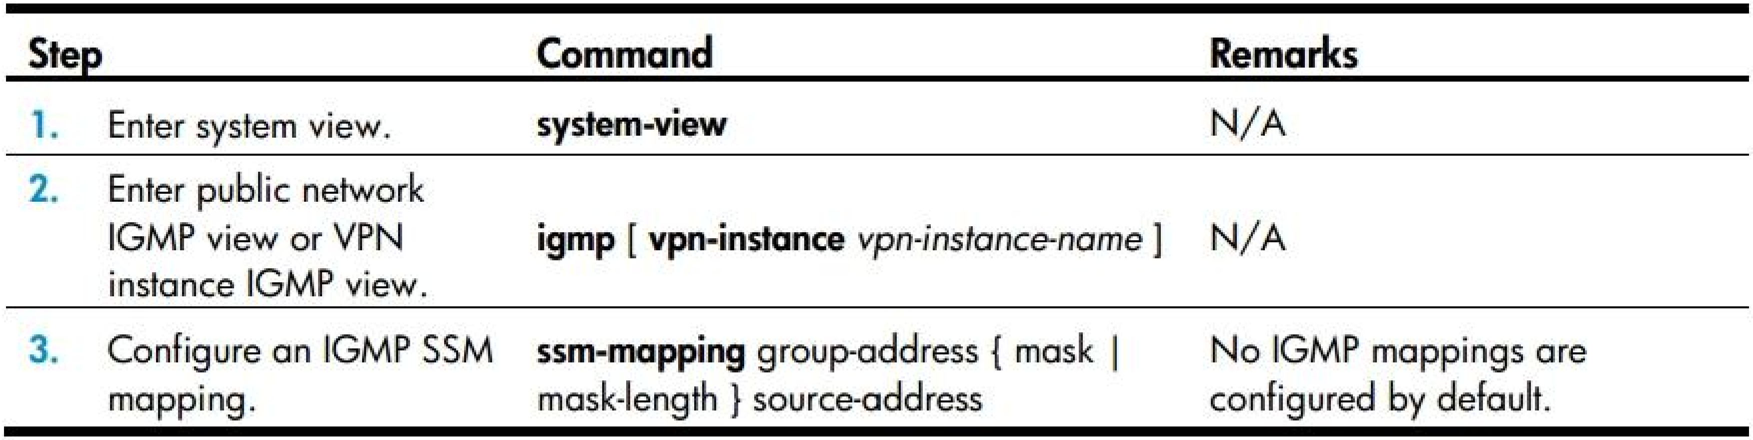

Step		Command	Remarks
1.	Enter system view.	**system-view**	N/A
2.	Enter public network IGMP view or VPN instance IGMP view.	**igmp** [**vpn-instance** *vpn-instance-name*]	N/A
3.	Configure an IGMP SSM mapping.	**ssm-mapping** group-address { mask | mask-length } source-address	No IGMP mappings are configured by default.

QUESTION 11
View the exhibit. PIM-SM is configured on each interface in the network and IGMP is enabled on VLAN 3 and VLAN 4, as shown in Exhibit 1. Each switch has a rendezvous point (RP) set and valid unicast routes, as shown in Exhibit 2. The network administrator configures an endpoint connected to VLAN4 to begin listening on 239.255.1.1 to verify if multicast routing configuration is established correctly. The network is not live, meaning no other hosts are sending IGMP requests. The network administrator then checks PIM routing entries on the switches to verify if the system is functioning correctly. On which switches should (*, 239.255.1.1) entries appear if the system is functioning correctly?

Exhibit 1

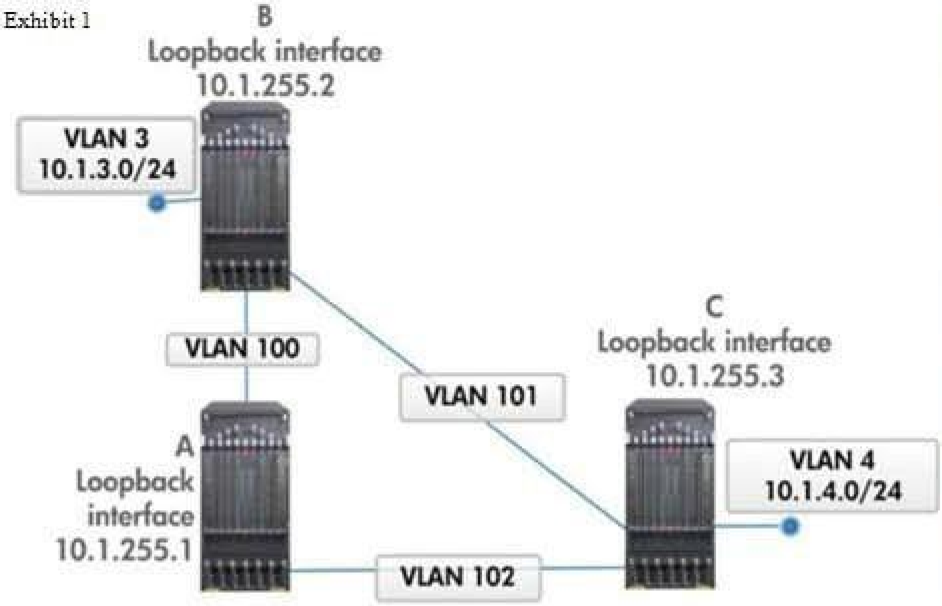

Exhibit 2

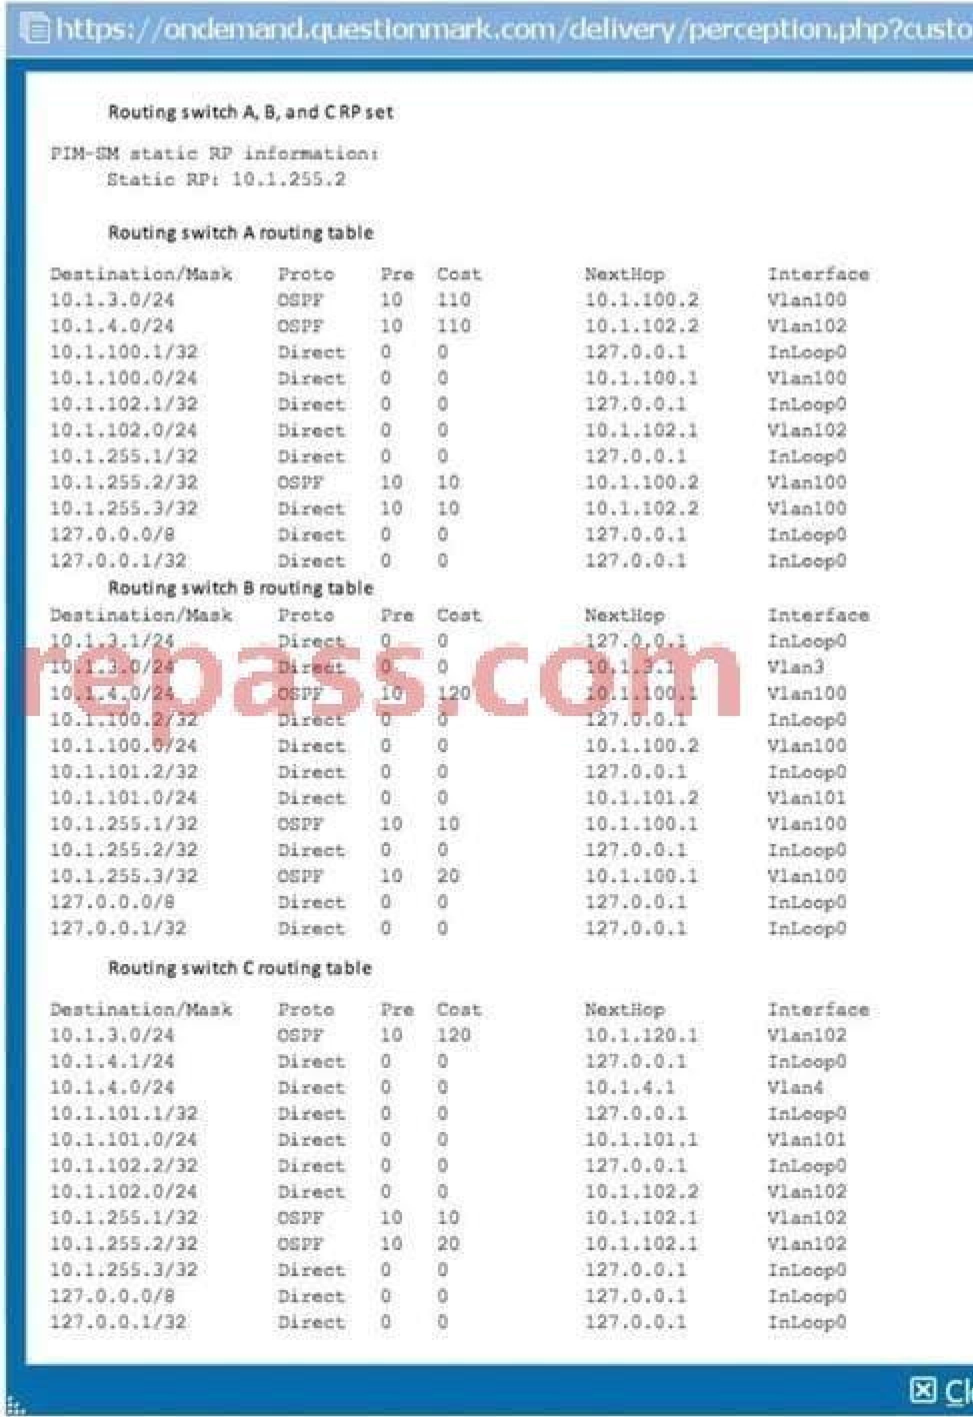
https://ondemand.questionmark.com/delivery/perception.php?custo
☒ Cl

Routing switch A, B, and C RP set

```
PIM-SM static RP information:
     Static RP: 10.1.255.2
```

Routing switch A routing table

```
Destination/Mask    Proto   Pre  Cost        NextHop       Interface
10.1.3.0/24         OSPF    10   110         10.1.100.2    Vlan100
10.1.4.0/24         OSPF    10   110         10.1.102.2    Vlan102
10.1.100.1/32       Direct  0    0           127.0.0.1     InLoop0
10.1.100.0/24       Direct  0    0           10.1.100.1    Vlan100
10.1.102.1/32       Direct  0    0           127.0.0.1     InLoop0
10.1.102.0/24       Direct  0    0           10.1.102.1    Vlan102
10.1.255.1/32       Direct  0    0           127.0.0.1     InLoop0
10.1.255.2/32       OSPF    10   10          10.1.100.2    Vlan100
10.1.255.3/32       Direct  10   10          10.1.102.2    Vlan100
127.0.0.0/8         Direct  0    0           127.0.0.1     InLoop0
127.0.0.1/32        Direct  0    0           127.0.0.1     InLoop0
```

Routing switch B routing table

```
Destination/Mask    Proto   Pre  Cost        NextHop       Interface
10.1.3.1/24         Direct  0    0           127.0.0.1     InLoop0
10.1.3.0/24         Direct  0    0           10.1.3.1      Vlan3
10.1.4.0/24         OSPF    10   120         10.1.100.1    Vlan100
10.1.100.2/32       Direct  0    0           127.0.0.1     InLoop0
10.1.100.0/24       Direct  0    0           10.1.100.2    Vlan100
10.1.101.2/32       Direct  0    0           127.0.0.1     InLoop0
10.1.101.0/24       Direct  0    0           10.1.101.2    Vlan101
10.1.255.1/32       OSPF    10   10          10.1.100.1    Vlan100
10.1.255.2/32       Direct  0    0           127.0.0.1     InLoop0
10.1.255.3/32       OSPF    10   20          10.1.100.1    Vlan100
127.0.0.0/8         Direct  0    0           127.0.0.1     InLoop0
127.0.0.1/32        Direct  0    0           127.0.0.1     InLoop0
```

Routing switch C routing table

```
Destination/Mask    Proto   Pre  Cost        NextHop       Interface
10.1.3.0/24         OSPF    10   120         10.1.120.1    Vlan102
10.1.4.1/24         Direct  0    0           127.0.0.1     InLoop0
10.1.4.0/24         Direct  0    0           10.1.4.1      Vlan4
10.1.101.1/32       Direct  0    0           127.0.0.1     InLoop0
10.1.101.0/24       Direct  0    0           10.1.101.1    Vlan101
10.1.102.2/32       Direct  0    0           127.0.0.1     InLoop0
10.1.102.0/24       Direct  0    0           10.1.102.2    Vlan102
10.1.255.1/32       OSPF    10   10          10.1.102.1    Vlan102
10.1.255.2/32       OSPF    10   20          10.1.102.1    Vlan102
10.1.255.3/32       Direct  0    0           127.0.0.1     InLoop0
127.0.0.0/8         Direct  0    0           127.0.0.1     InLoop0
127.0.0.1/32        Direct  0    0           127.0.0.1     InLoop0
```

A. switch B only
B. switch C only
C. switches B and C only
D. switches A, B, and C

Correct Answer: D
Section: Multicast
Explanation

Explanation/Reference:
In accordance with the Exhibit 2, a point of rendezvous (RP) is 10.1.255.2. According to the routing table of Switches, a message will be passed along the way: C -> A -> B. Thus, entry (*, G) will be created on the switches A, B and C.

HP 10500 Switch Series IP Multicast Configuration Guide

RPT building

Figure 42 RPT building in a PIM-SM domain

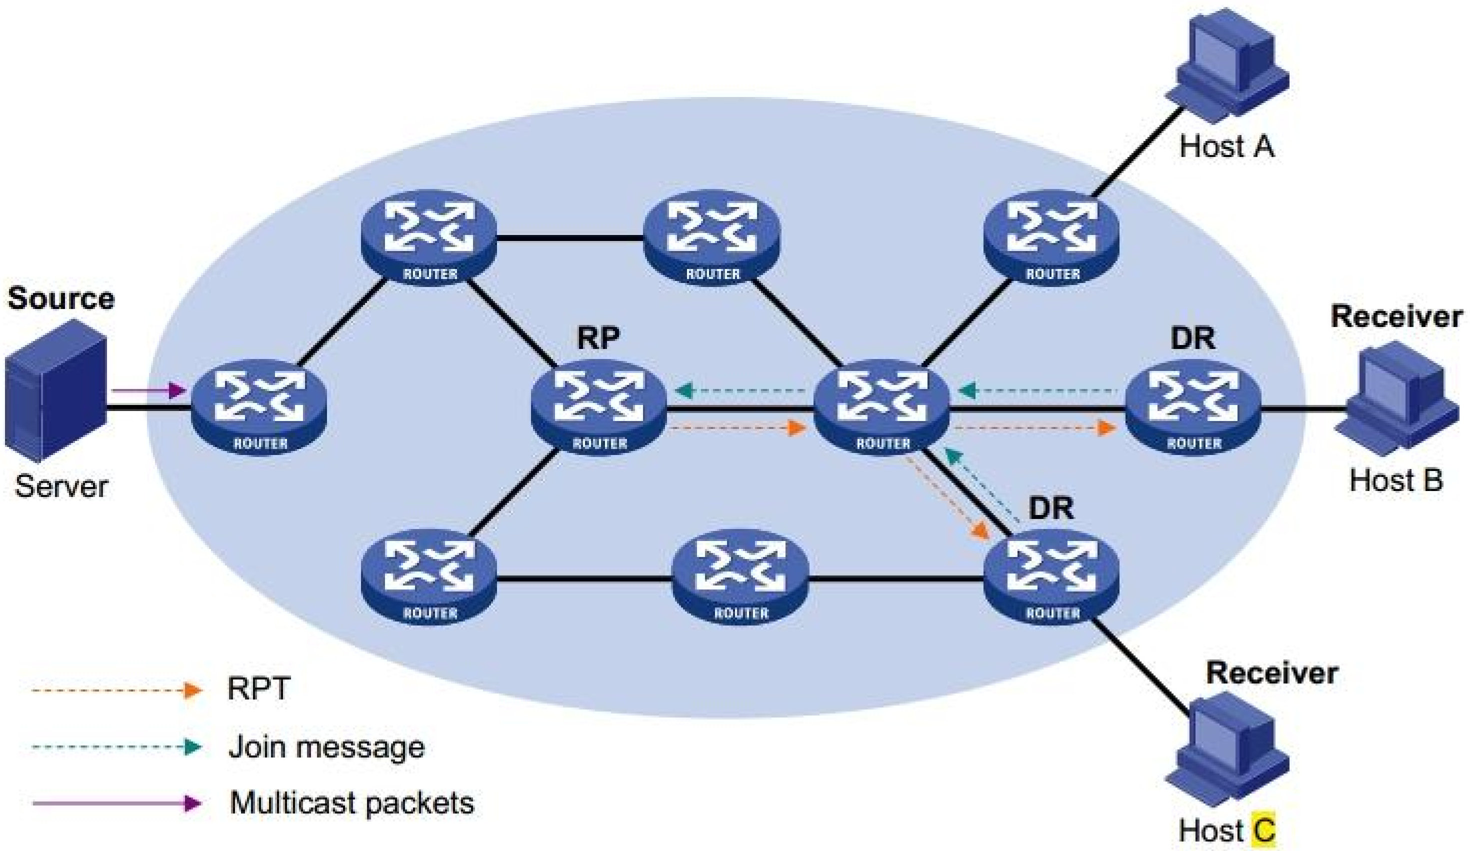

As shown in Figure 42, the process of building an RPT is as follows:

1. When a receiver joins multicast group G, it uses an IGMP message to inform the directly connec DR.
2. After getting the receiver information, the DR sends a join message, which is forwarded, hop by hop, to the RP that corresponds to the multicast group.
3. The routers along the path from the DR to the RP form an RPT branch. Each router on this branch

QUESTION 12
View the exhibit. A host connected to routing switch A joins group 239.0.0.1. Which switch does routing switch A select as the RP?

Exhibit 1

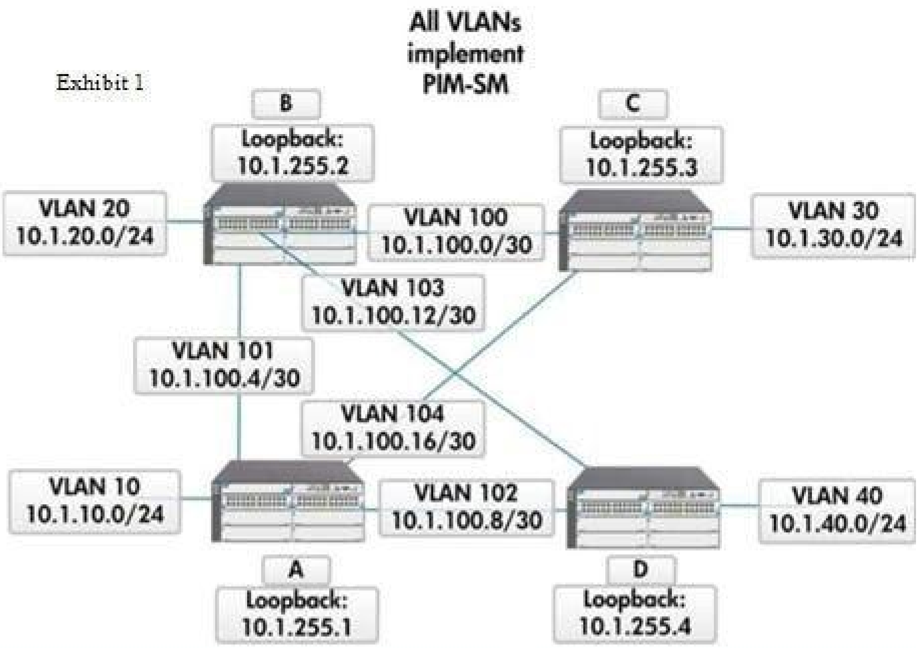

Exhibit 2

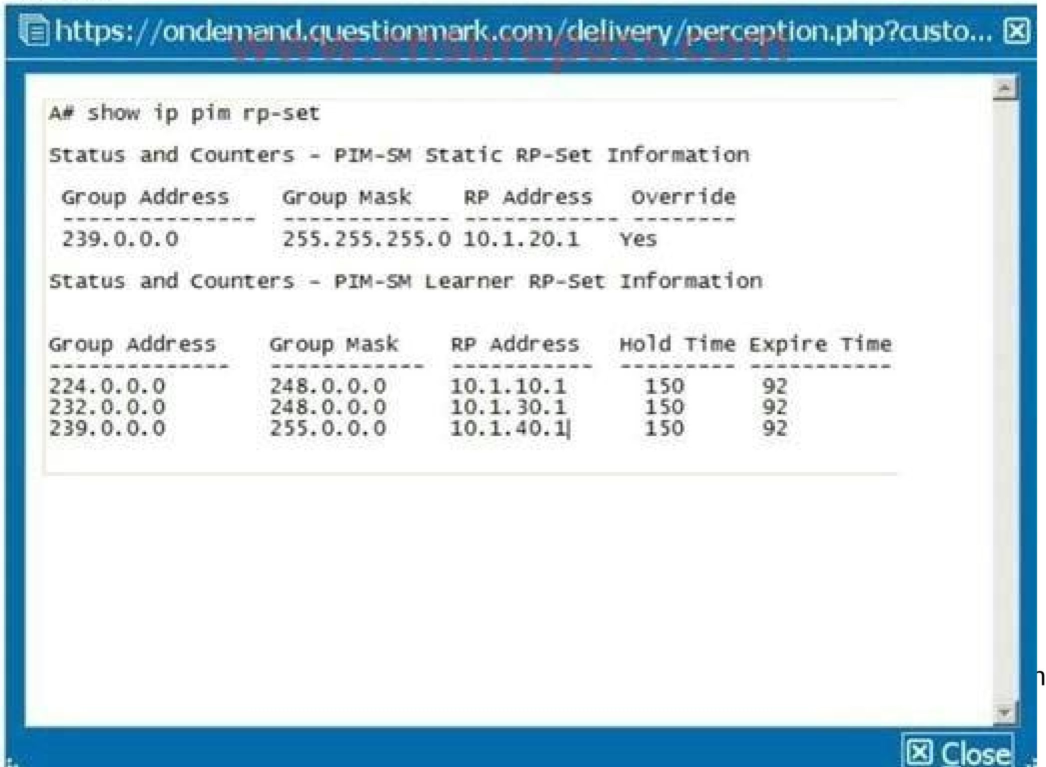
https://ondemand.questionmark.com/delivery/perception.php?custo...
Close

```
A# show ip pim rp-set

Status and Counters - PIM-SM Static RP-Set Information

 Group Address    Group Mask      RP Address   Override
 ---------------  --------------  ------------ --------
 239.0.0.0        255.255.255.0 10.1.20.1    Yes

Status and Counters - PIM-SM Learner RP-Set Information


Group Address     Group Mask      RP Address    Hold Time Expire Time
--------------    ------------    -----------   --------- -----------
224.0.0.0         248.0.0.0       10.1.10.1       150       92
232.0.0.0         248.0.0.0       10.1.30.1       150       92
239.0.0.0         255.0.0.0       10.1.40.1       150       92
```

A. Routing switch D
B. Routing switch C
C. Routing switch B
D. Itself

Correct Answer: C
Section: Multicast
Explanation

Explanation/Reference:
See RP-address column in the output.

QUESTION 13
A network uses PIM-DM to route multicast traffic. The network administrator has noticed bursts of congestion related to periodic floods of multicast traffic. How can the administrator eliminate these periodic bursts?

A. Leave PIM-DM enabled on the interfaces that face receivers, but enable PIM-SM throughout the core.
B. Decrease the PIM hello timer on VLAN interfaces on which PIM-DM is enabled.
C. Decrease the IGMP robustness variable on VLAN interfaces that connect to receivers.
D. Enable the state refresh feature on the VLAN interfaces on which PIM-DM is enabled.

Correct Answer: D
Section: Multicast
Explanation

Explanation/Reference:
HP 10500 Switch Series IP Multicast Configuration Guide

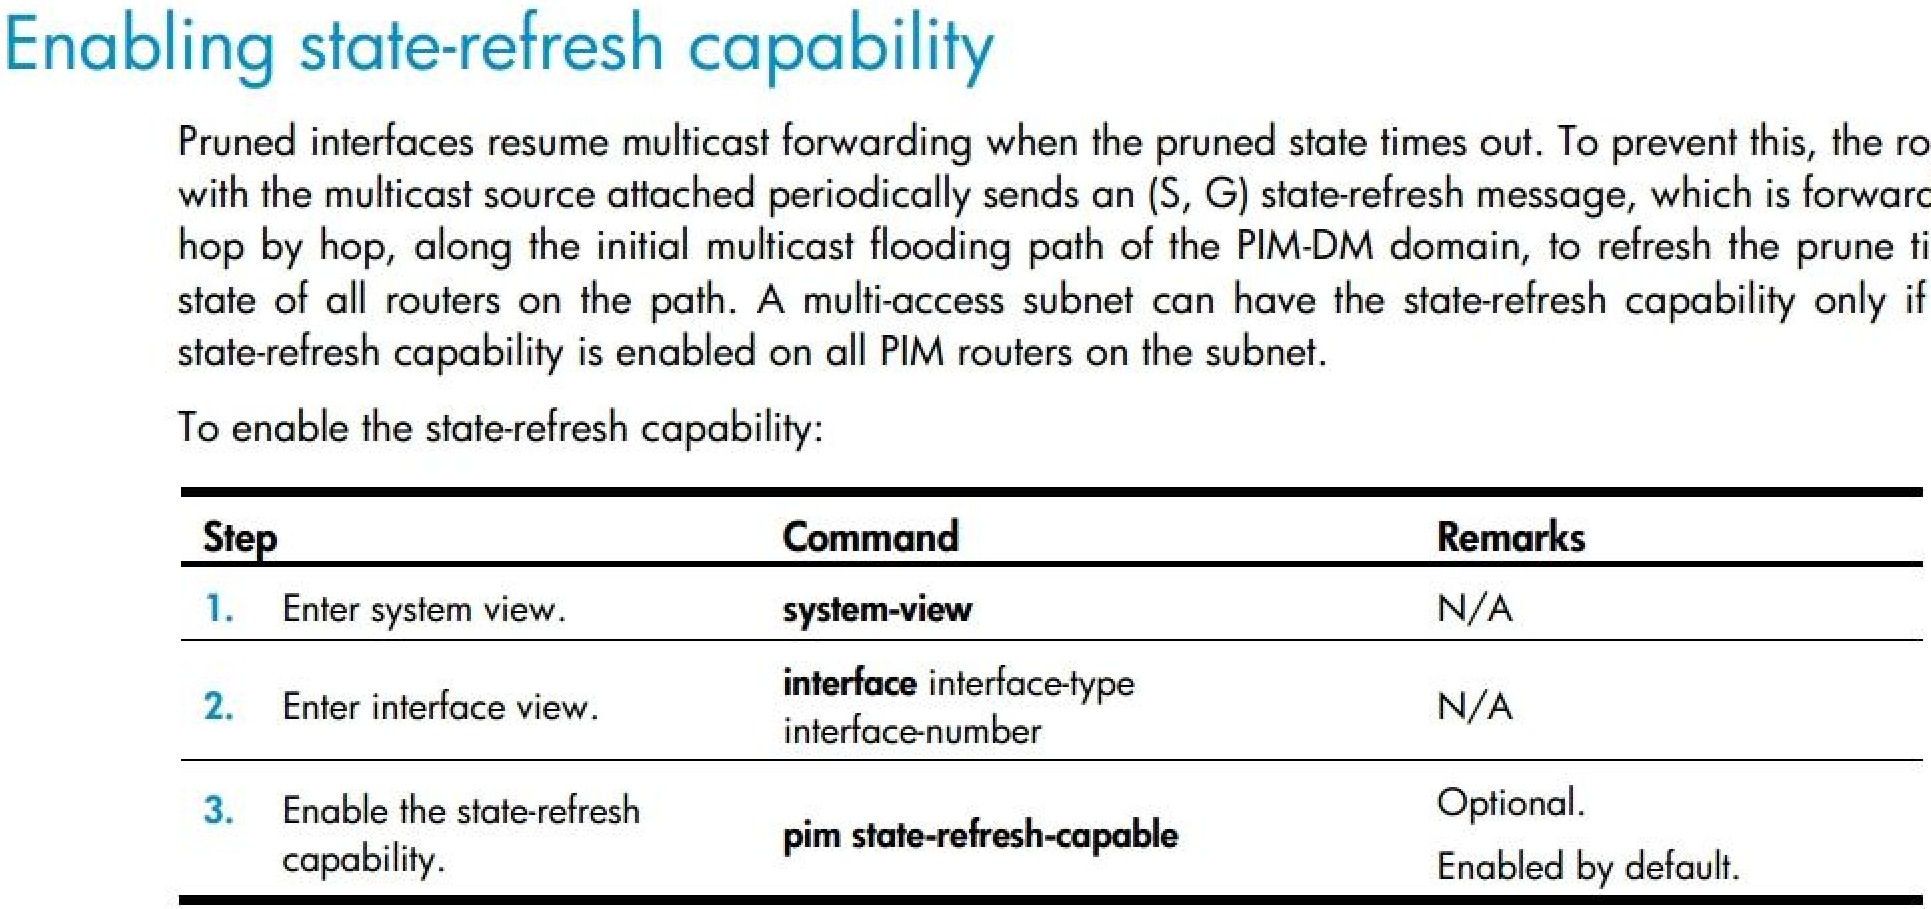

Enabling state-refresh capability

Pruned interfaces resume multicast forwarding when the pruned state times out. To prevent this, the ro with the multicast source attached periodically sends an (S, G) state-refresh message, which is forward hop by hop, along the initial multicast flooding path of the PIM-DM domain, to refresh the prune ti state of all routers on the path. A multi-access subnet can have the state-refresh capability only if state-refresh capability is enabled on all PIM routers on the subnet.

To enable the state-refresh capability:

Step		Command	Remarks
1.	Enter system view.	**system-view**	N/A
2.	Enter interface view.	**interface** *interface-type interface-number*	N/A
3.	Enable the state-refresh capability.	**pim state-refresh-capable**	Optional. Enabled by default.

QUESTION 14
View the exhibit. All switches shown in Exhibit 1 are HP Provision ASIC switches, which are implementing IGMP and multicast routing with PIM. Routing switch B has just received the message shown in Exhibit 2 from routing switch A. What can the network administrator assume?

Exhibit 1

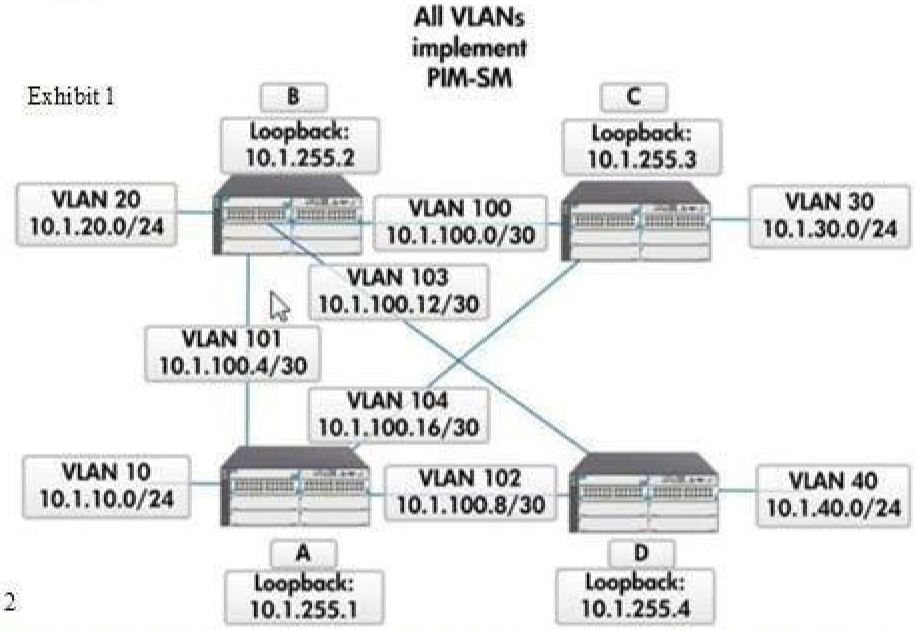

Exhibit 2

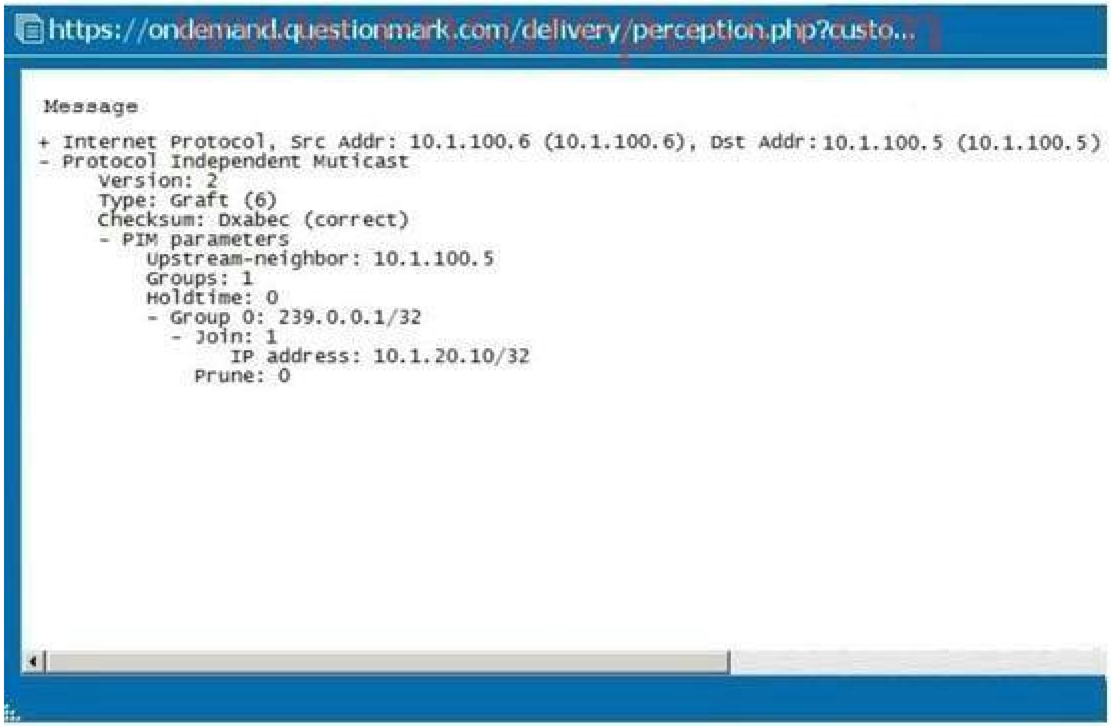

https://ondemand.questionmark.com/delivery/perception.php?custo...

```
Message

+ Internet Protocol, Src Addr: 10.1.100.6 (10.1.100.6), Dst Addr:10.1.100.5 (10.1.100.5)
- Protocol Independent Muticast
      Version: 2
      Type: Graft (6)
      Checksum: Dxabec (correct)
    - PIM parameters
          Upstream-neighbor: 10.1.100.5
          Groups: 1
          Holdtime: 0
        - Group 0: 239.0.0.1/32
          - Join: 1
                IP address: 10.1.20.10/32
            Prune: 0
```

A. Routing switch B is responding to routing switch A's state refresh inquiry.
B. A host connected to routing switch B has joined the 239.0.0.1 group, and routing switch B wants to join the RP tree for this group.
C. Previously, routing switch B did not require 239.0.0.1 multicasts, but it now does due to a change in topology or IGMP reports.
D. A new multicast source for 239.0.0.1 has begun to transmit in a VLAN connected to routing switch B.

Correct Answer: B
Section: Multicast
Explanation

Explanation/Reference:

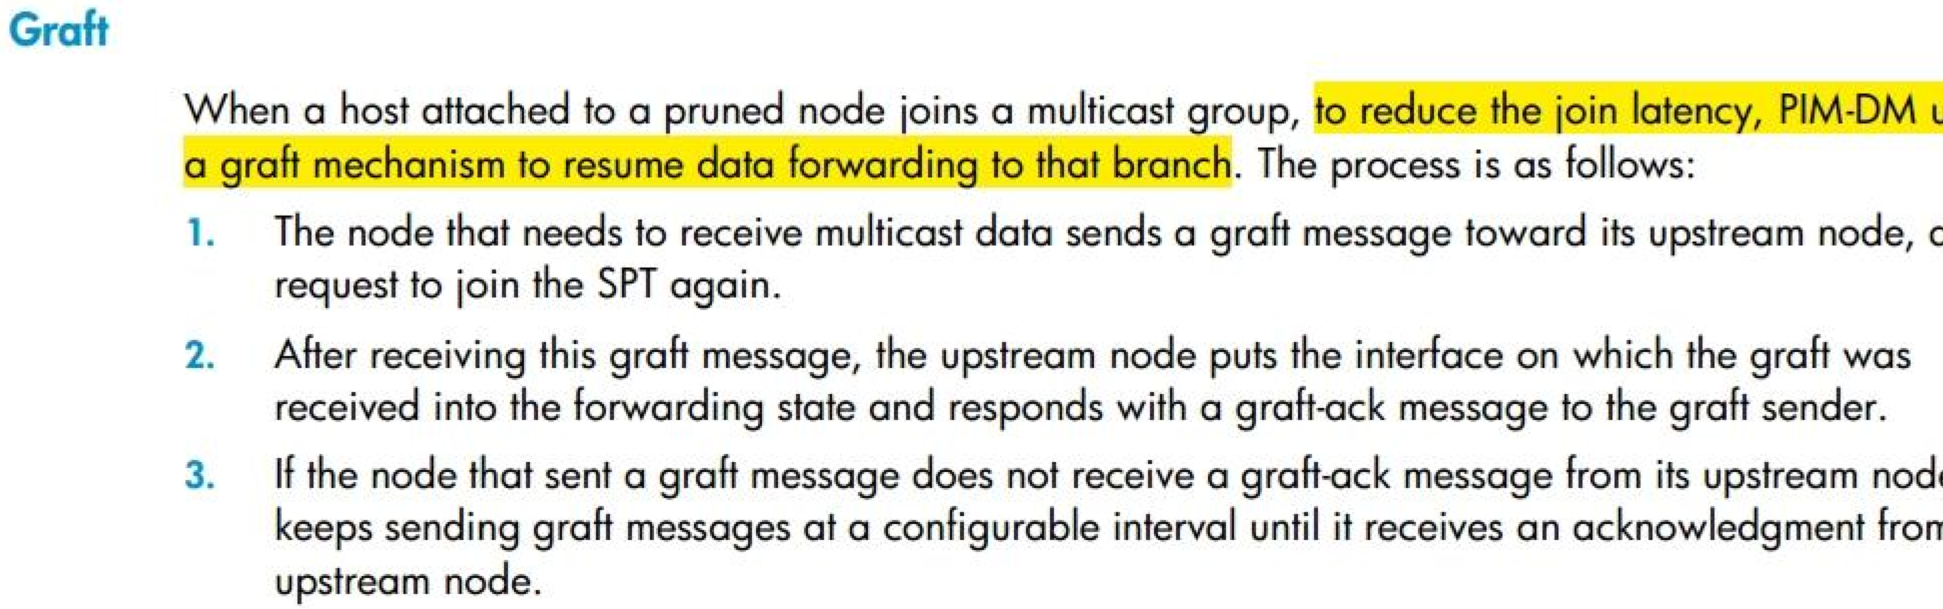

Graft

When a host attached to a pruned node joins a multicast group, to reduce the join latency, PIM-DM u a graft mechanism to resume data forwarding to that branch. The process is as follows:

1. The node that needs to receive multicast data sends a graft message toward its upstream node, c request to join the SPT again.
2. After receiving this graft message, the upstream node puts the interface on which the graft was received into the forwarding state and responds with a graft-ack message to the graft sender.
3. If the node that sent a graft message does not receive a graft-ack message from its upstream nod keeps sending graft messages at a configurable interval until it receives an acknowledgment fron upstream node.

QUESTION 15
View the exhibit. An HP 5400 zl switch implements the configuration shown in the exhibit. However, when the network administrator tests the configuration, endpoints in VLAN 10 do not obtain global IPv6 addresses. What might be causing this issue?

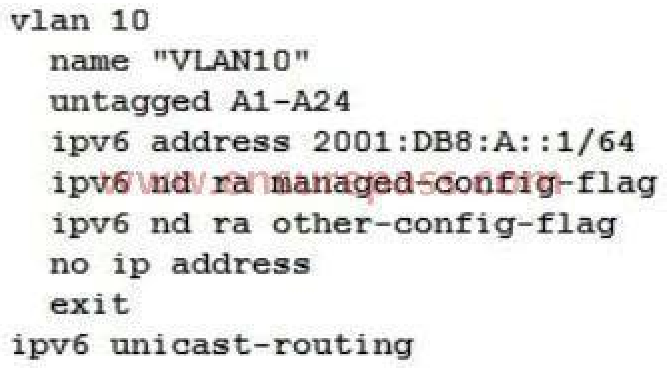

```
vlan 10
  name "VLAN10"
  untagged A1-A24
  ipv6 address 2001:DB8:A::1/64
  ipv6 nd ra managed-config-flag
  ipv6 nd ra other-config-flag
  no ip address
  exit
ipv6 unicast-routing
```

A. The switch is suppressing router advertisement (RA) suppression on the VLAN.
B. The switch has not been enabled for IPv6 globally.
C. The switch is not implementing DHCPv6 relay.
D. The switch is not configured to advertise the correct prefix.

Correct Answer: C
Section: IPv6
Explanation

Explanation/Reference:
IPv6 hosts can configure themselves automatically when connected to an IPv6 network using the Neighbor Discovery Protocol via Internet Control Message Protocol version 6 (ICMPv6) router discovery messages. When first connected to a network, a host sends a link-local router solicitation multicast request for its configuration parameters; routers respond to such a request with a router advertisement packet that contains Internet Layer configuration parameters.
If IPv6 stateless address autoconfiguration is unsuitable for an application, a network may use stateful configuration with the Dynamic Host Configuration Protocol version 6 (DHCPv6) or hosts may be configured manually using static methods.

In question apparently assumed to that the DHCPv6 server is on a different network, and therefore the routing and DHCP-relay required!

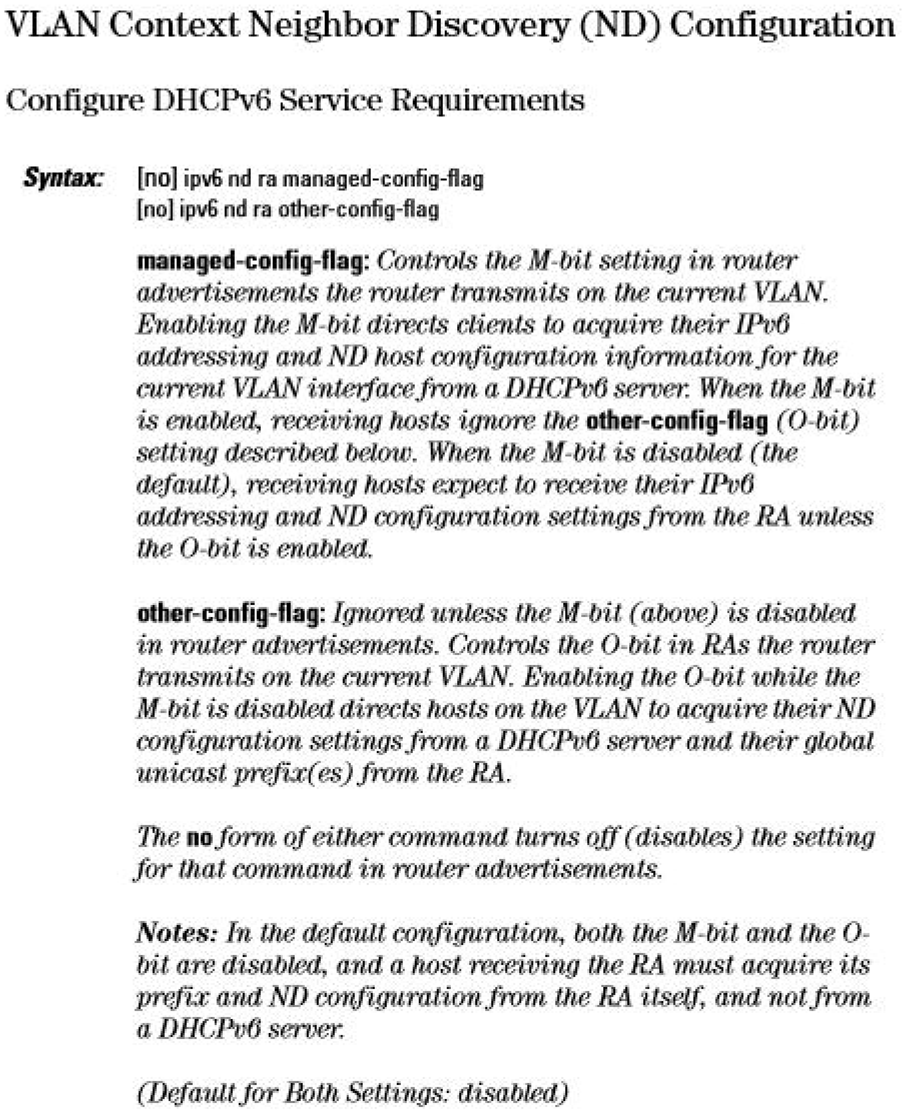

VLAN Context Neighbor Discovery (ND) Configuration

Configure DHCPv6 Service Requirements

Syntax: [no] ipv6 nd ra managed-config-flag
[no] ipv6 nd ra other-config-flag

managed-config-flag: *Controls the M-bit setting in router advertisements the router transmits on the current VLAN. Enabling the M-bit directs clients to acquire their IPv6 addressing and ND host configuration information for the current VLAN interface from a DHCPv6 server. When the M-bit is enabled, receiving hosts ignore the* **other-config-flag** *(O-bit) setting described below. When the M-bit is disabled (the default), receiving hosts expect to receive their IPv6 addressing and ND configuration settings from the RA unless the O-bit is enabled.*

other-config-flag: *Ignored unless the M-bit (above) is disabled in router advertisements. Controls the O-bit in RAs the router transmits on the current VLAN. Enabling the O-bit while the M-bit is disabled directs hosts on the VLAN to acquire their ND configuration settings from a DHCPv6 server and their global unicast prefix(es) from the RA.*

The **no** *form of either command turns off (disables) the setting for that command in router advertisements.*

Notes: *In the default configuration, both the M-bit and the O-bit are disabled, and a host receiving the RA must acquire its prefix and ND configuration from the RA itself, and not from a DHCPv6 server.*

(Default for Both Settings: disabled)

QUESTION 16
View the exhibit. A site features both IPv6 and IPv4 endpoints and IPv4 and IPv6 capable servers in the data center, as shown in the exhibit. How can the network administrator configure the HP 8200 zl switch at the core to work in this environment?

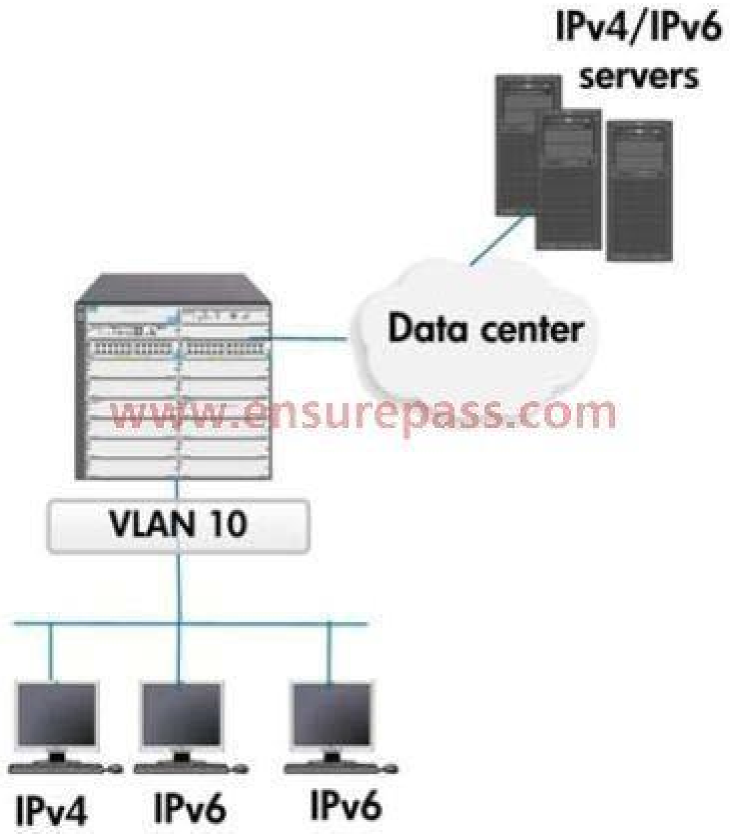

A. Configure IPv4 and IPv6 addresses on VLAN 10 and set up IPv4 and IPv6 routing solutions
B. Configure a static IPv6 over IPv4 tunnel between the campus LAN and the data center
C. Configure the switch to communicate with a Teredo broker on VLAN 10
D. Configure ISATAP tunnels in VLAN 10 and set up IPv4 and IPv6 routing solutions

Correct Answer: D
Section: IPv6
Explanation

Explanation/Reference:
Here the main problem provide interact of IPv4 devices to with IPv6 devices on the same local network. Tunneling can not be applied. Therefore required address translation, which performs ISATAP.

QUESTION 17
Two HP 7500 switches, switch A and switch B, are implementing VRRP in VLAN 2 (VRID 2) and in VLAN 3 (VRID 3). Switch A is currently the master in both VRIDs. The network administrator enables load balancing on both VRIDs on both switches. Which statement describes an effect of these commands?

A. Switch A becomes master in one VRID, and switch B becomes master in the other VRID.
B. Switch B requests and receives its own virtual MAC address in both VRIDs from switch A.
C. Switch A and switch B respond to ARP requests in both VLANs.
D. Switch A and switch B send out standard VRRP advertisements in both VLANs.

Correct Answer: B
Section: High Availability
Explanation

Explanation/Reference:
Assumes that the system administrator perform the command **vrrp mode load-balance** on each switch.

VRRP Load Balancing Mode

Overview

When VRRP works in standard protocol mode, only the master can forward packets and the backups are in the state of listening. You can create multiple VRRP groups to share the load among multiple routers, but hosts on the LAN need to be configured with different gateways, thus making the configuration complicated.

In load balancing mode, VRRP provides load balancing in addition to virtual gateway redundancy by mapping a virtual IP address to multiple virtual MAC addresses to assign each router in a VRRP group one virtual MAC address. In this way, each router in this VRRP group can forward packets. In load balancing mode, you need to create only one VRRP group to balance load among multiple routers, instead of allowing one router to bear the load while other routers stay idle.

NOTE:

VRRP load balancing mode is based on VRRP standard protocol mode, so mechanisms, such as master election, preemption, and tracking functions, in the standard protocol mode are also supported in the load balancing mode. In addition, VRRP load balancing mode has new mechanisms, which are introduced in the following sections.

Assigning Virtual MAC Addresses

When VRRP works in load balancing mode, the master assigns virtual MAC addresses to routers in the VRRP group and answers the ARP requests (for the IPv4 network) or ND requests (for the IPv6 network) from different hosts. The backup routers, however, do not answer the ARP requests (for the IPv4 network) or ND requests (for the IPv6 network) from the hosts.

Figure 7 Allocating virtual MAC addresses

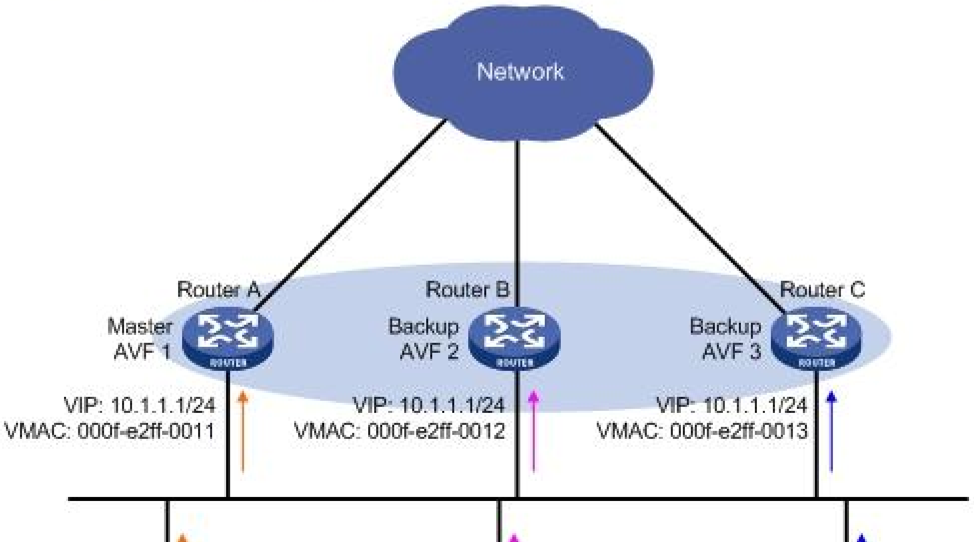

QUESTION 18
Which type of IPv6 over IPv4 tunnel protocol on an HP 12500 switch requires that you specify a destination?

A. IPv6 over IPv4
B. ISATAP
C. IPv6 over IPv4 auto (or IPv4-compatible)
D. 6 to 4

Correct Answer: A
Section: IPv6
Explanation

Explanation/Reference:

Table 11 IPv6 over IPv4 tunnel modes and key parameters

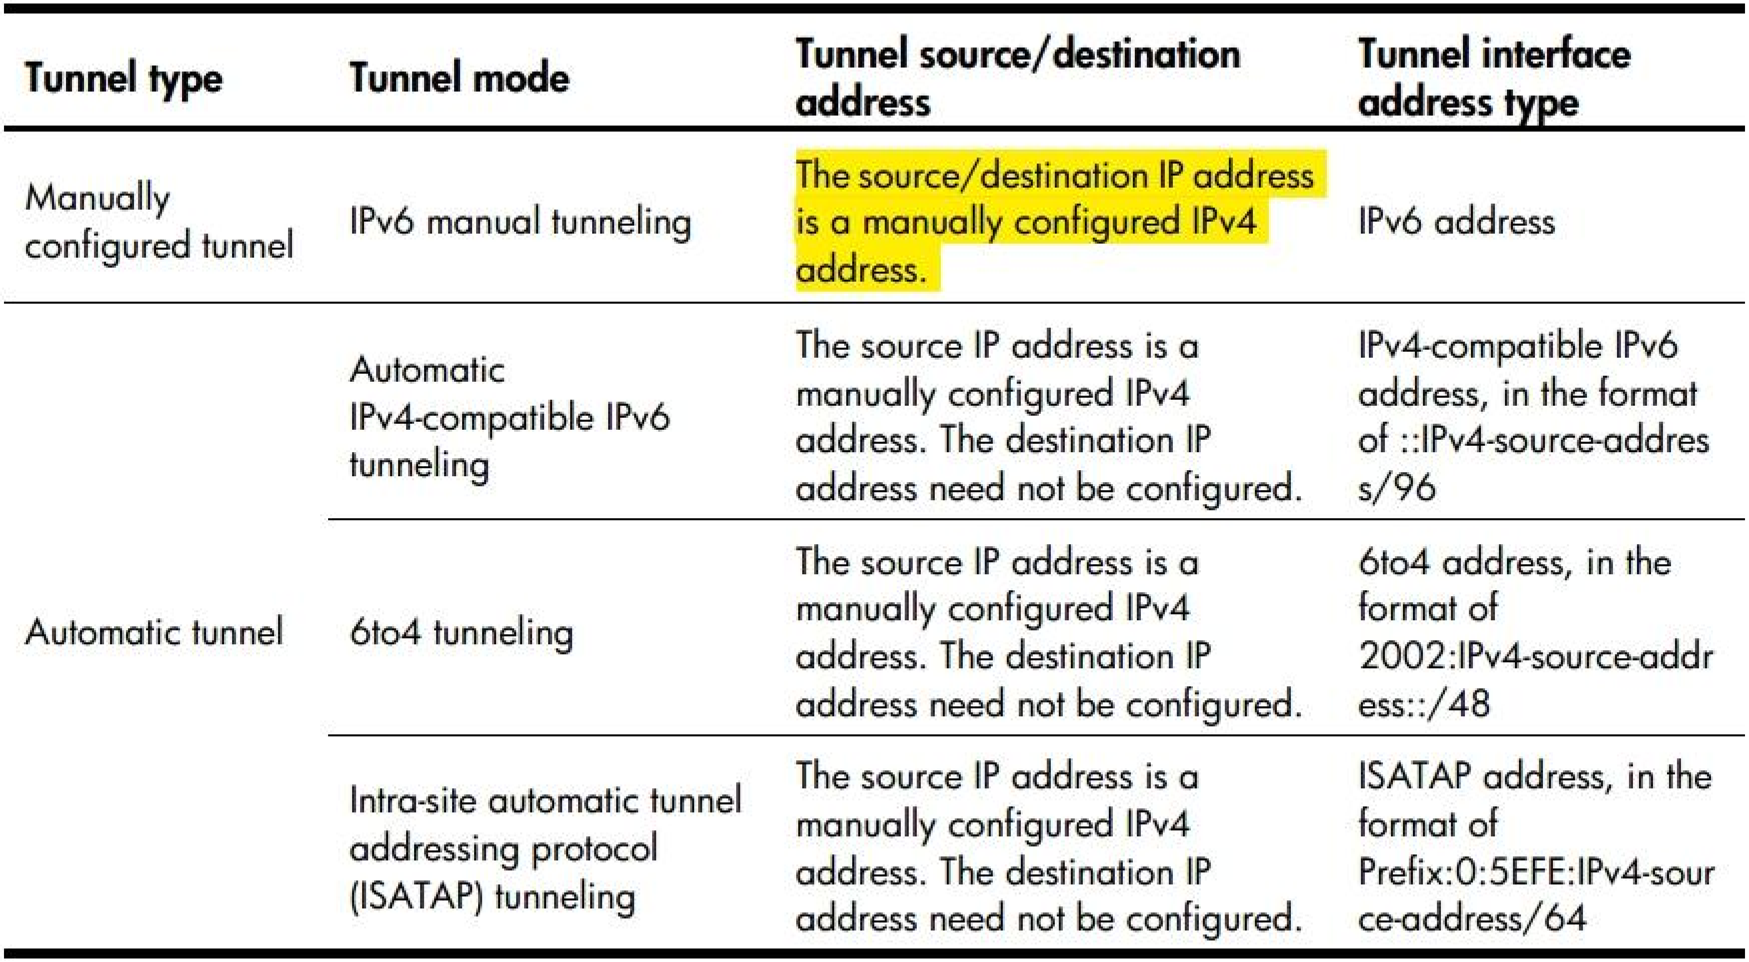

Tunnel type	Tunnel mode	Tunnel source/destination address	Tunnel interface address type
Manually configured tunnel	IPv6 manual tunneling	The source/destination IP address is a manually configured IPv4 address.	IPv6 address
Automatic tunnel	Automatic IPv4-compatible IPv6 tunneling	The source IP address is a manually configured IPv4 address. The destination IP address need not be configured.	IPv4-compatible IPv6 address, in the format of ::IPv4-source-address/96
	6to4 tunneling	The source IP address is a manually configured IPv4 address. The destination IP address need not be configured.	6to4 address, in the format of 2002:IPv4-source-address::/48
	Intra-site automatic tunnel addressing protocol (ISATAP) tunneling	The source IP address is a manually configured IPv4 address. The destination IP address need not be configured.	ISATAP address, in the format of Prefix:0:5EFE:IPv4-source-address/64

QUESTION 19
A network administrator enabled IPv6 on an HP 7500 switch and configured VLAN interface 10 with a link-local and global IPv6 address 2001:DB8:1:10::1/64. The network administrator now wants IPv6-capable endpoints connected to VLAN 10 to receive IPv6 addresses using stateless autoconfiguration. Which step must the network administrator complete on the 7500 switch?

A. Set the managed configuration flag in RA messages.

B. Configure DHCPv6 relay on VLAN interface 10.
C. Re-enable RA messages on VLAN interface 10.
D. Configure the VLAN 10 interface to advertise prefix 2001:DB8:1:10::/64.

Correct Answer: C
Section: IPv6
Explanation

Explanation/Reference:
HP 7500 IP Service Configuration guide

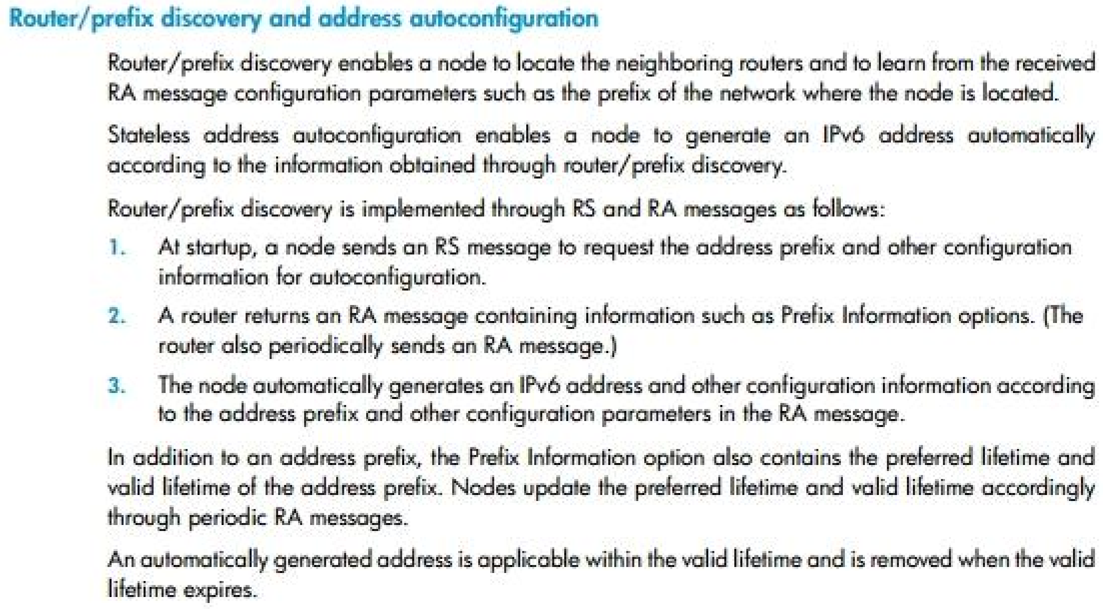

Router/prefix discovery and address autoconfiguration

Router/prefix discovery enables a node to locate the neighboring routers and to learn from the received RA message configuration parameters such as the prefix of the network where the node is located.

Stateless address autoconfiguration enables a node to generate an IPv6 address automatically according to the information obtained through router/prefix discovery.

Router/prefix discovery is implemented through RS and RA messages as follows:

1. At startup, a node sends an RS message to request the address prefix and other configuration information for autoconfiguration.
2. A router returns an RA message containing information such as Prefix Information options. (The router also periodically sends an RA message.)
3. The node automatically generates an IPv6 address and other configuration information according to the address prefix and other configuration parameters in the RA message.

In addition to an address prefix, the Prefix Information option also contains the preferred lifetime and valid lifetime of the address prefix. Nodes update the preferred lifetime and valid lifetime accordingly through periodic RA messages.

An automatically generated address is applicable within the valid lifetime and is removed when the valid lifetime expires.

...

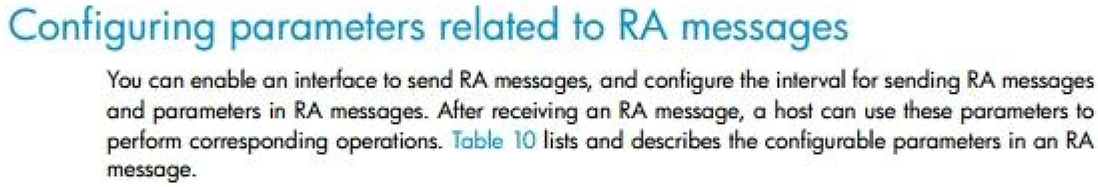

Configuring parameters related to RA messages

You can enable an interface to send RA messages, and configure the interval for sending RA messages and parameters in RA messages. After receiving an RA message, a host can use these parameters to perform corresponding operations. Table 10 lists and describes the configurable parameters in an RA message.

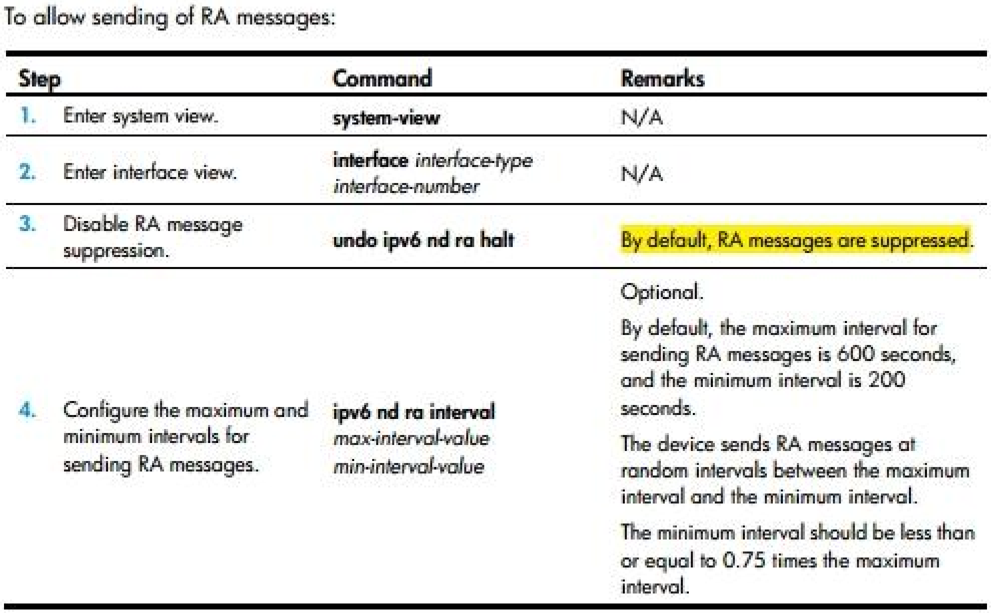

To allow sending of RA messages:

Step		Command	Remarks
1.	Enter system view.	**system-view**	N/A
2.	Enter interface view.	**interface** *interface-type interface-number*	N/A
3.	Disable RA message suppression.	**undo ipv6 nd ra halt**	By default, RA messages are suppressed.
4.	Configure the maximum and minimum intervals for sending RA messages.	**ipv6 nd ra interval** *max-interval-value min-interval-value*	Optional. By default, the maximum interval for sending RA messages is 600 seconds, and the minimum interval is 200 seconds. The device sends RA messages at random intervals between the maximum interval and the minimum interval. The minimum interval should be less than or equal to 0.75 times the maximum interval.

QUESTION 20

View the Exhibit. A network administrator configures VRRP on two HP 5800 switches, as shown in the exhibit. VRRP settings not shown in the running-config are at their default settings. Which statement accurately describes the VRRP group?

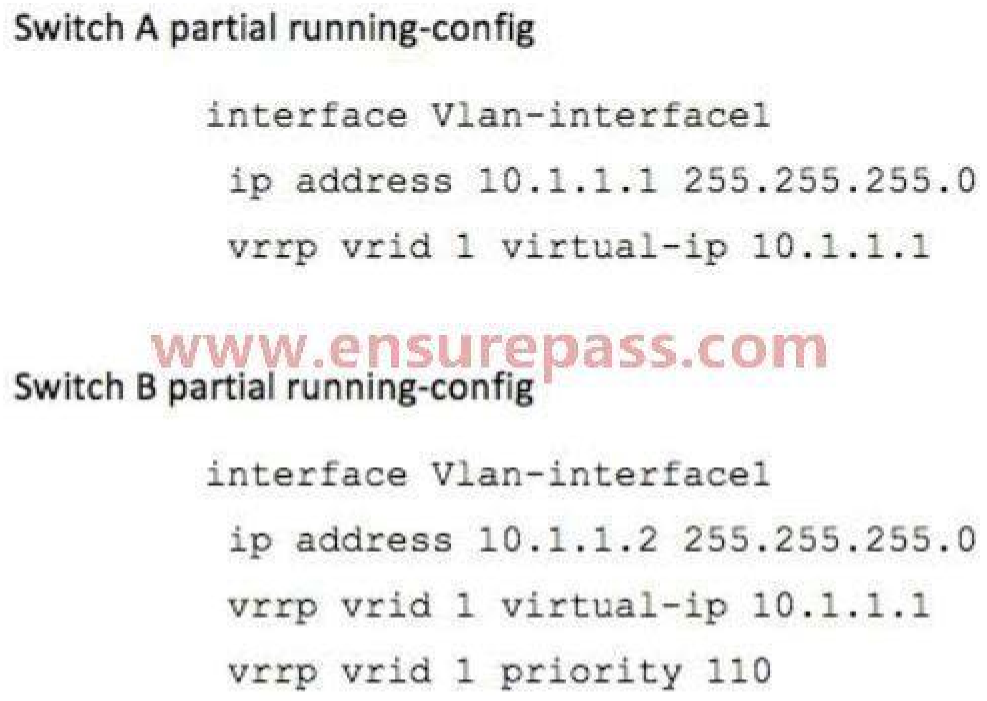

Switch A partial running-config

```
interface Vlan-interface1
 ip address 10.1.1.1 255.255.255.0
 vrrp vrid 1 virtual-ip 10.1.1.1
```

Switch B partial running-config

```
interface Vlan-interface1
 ip address 10.1.1.2 255.255.255.0
 vrrp vrid 1 virtual-ip 10.1.1.1
 vrrp vrid 1 priority 110
```

A. Switch A has effective priority 255 and always acts as master when it is up.
B. Switch B acts as Master when it is up. However, if switch B fails and comes back up, switch A remains master.
C. Switch B has effective priority 110 and always acts as master when it is up.
D. Switch A acts as master when it is up. However, if switch A fails and comes back up, switch B remains master.

Correct Answer: A
Section: High Availability
Explanation

Explanation/Reference:
Because Switch A is owner of the IP address of VRRP group, then his priority is 255. VRRP group on default working in preemption mode, therefore Switch A with highest priority is always act as master.

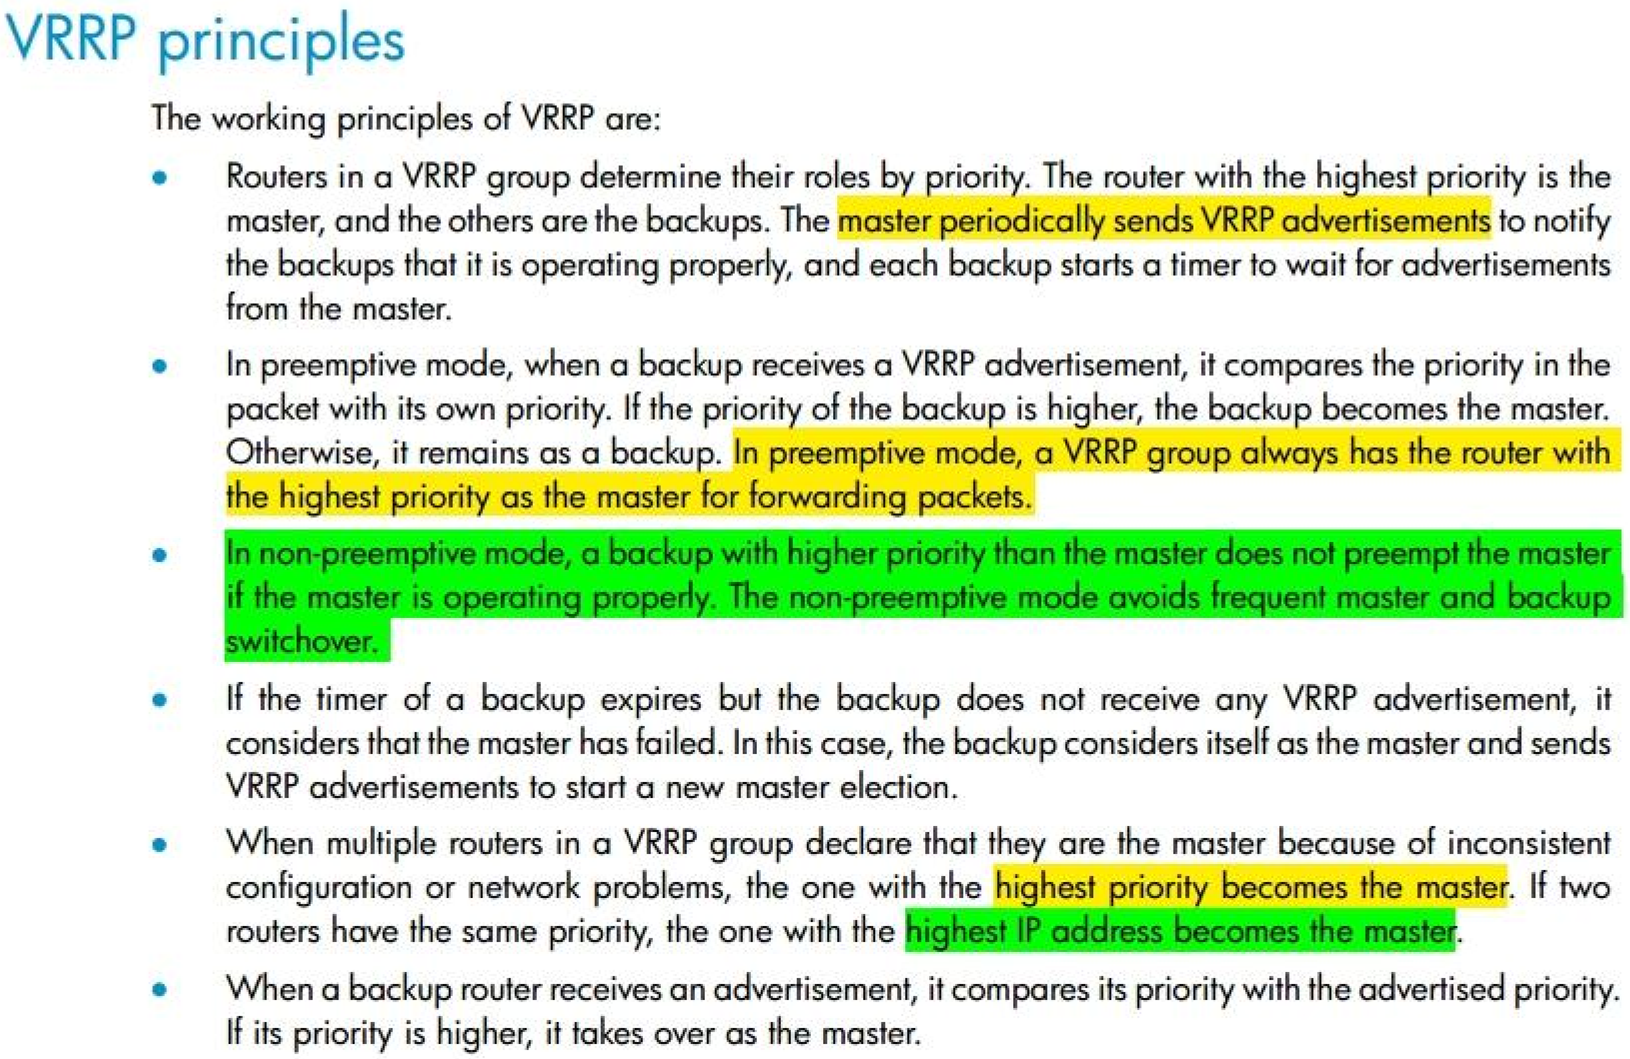

VRRP principles

The working principles of VRRP are:

- Routers in a VRRP group determine their roles by priority. The router with the highest priority is the master, and the others are the backups. The master periodically sends VRRP advertisements to notify the backups that it is operating properly, and each backup starts a timer to wait for advertisements from the master.
- In preemptive mode, when a backup receives a VRRP advertisement, it compares the priority in the packet with its own priority. If the priority of the backup is higher, the backup becomes the master. Otherwise, it remains as a backup. In preemptive mode, a VRRP group always has the router with the highest priority as the master for forwarding packets.
- In non-preemptive mode, a backup with higher priority than the master does not preempt the master if the master is operating properly. The non-preemptive mode avoids frequent master and backup switchover.
- If the timer of a backup expires but the backup does not receive any VRRP advertisement, it considers that the master has failed. In this case, the backup considers itself as the master and sends VRRP advertisements to start a new master election.
- When multiple routers in a VRRP group declare that they are the master because of inconsistent configuration or network problems, the one with the highest priority becomes the master. If two routers have the same priority, the one with the highest IP address becomes the master.
- When a backup router receives an advertisement, it compares its priority with the advertised priority. If its priority is higher, it takes over as the master.

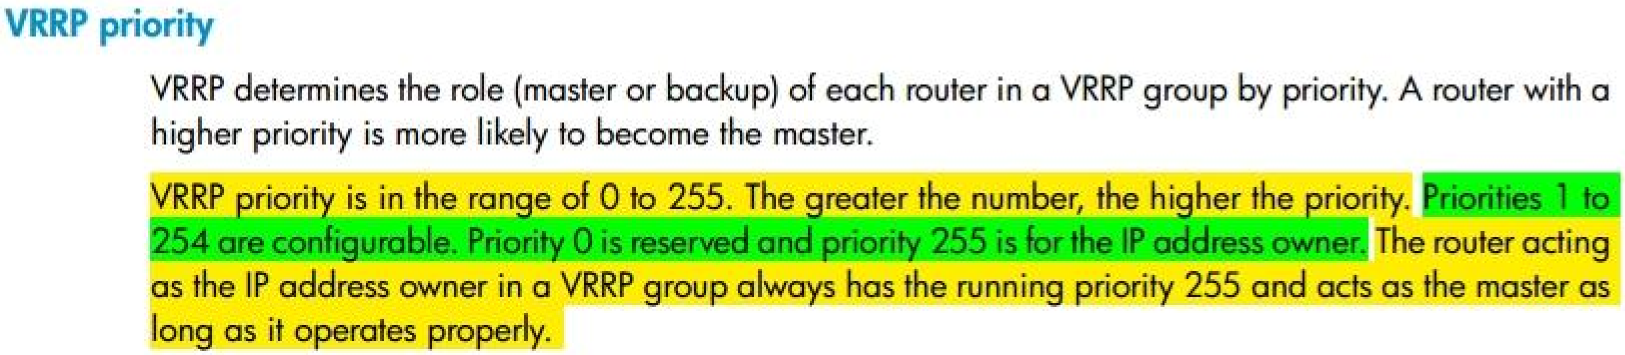

VRRP priority

VRRP determines the role (master or backup) of each router in a VRRP group by priority. A router with a higher priority is more likely to become the master.

VRRP priority is in the range of 0 to 255. The greater the number, the higher the priority. Priorities 1 to 254 are configurable. Priority 0 is reserved and priority 255 is for the IP address owner. The router acting as the IP address owner in a VRRP group always has the running priority 255 and acts as the master as long as it operates properly.

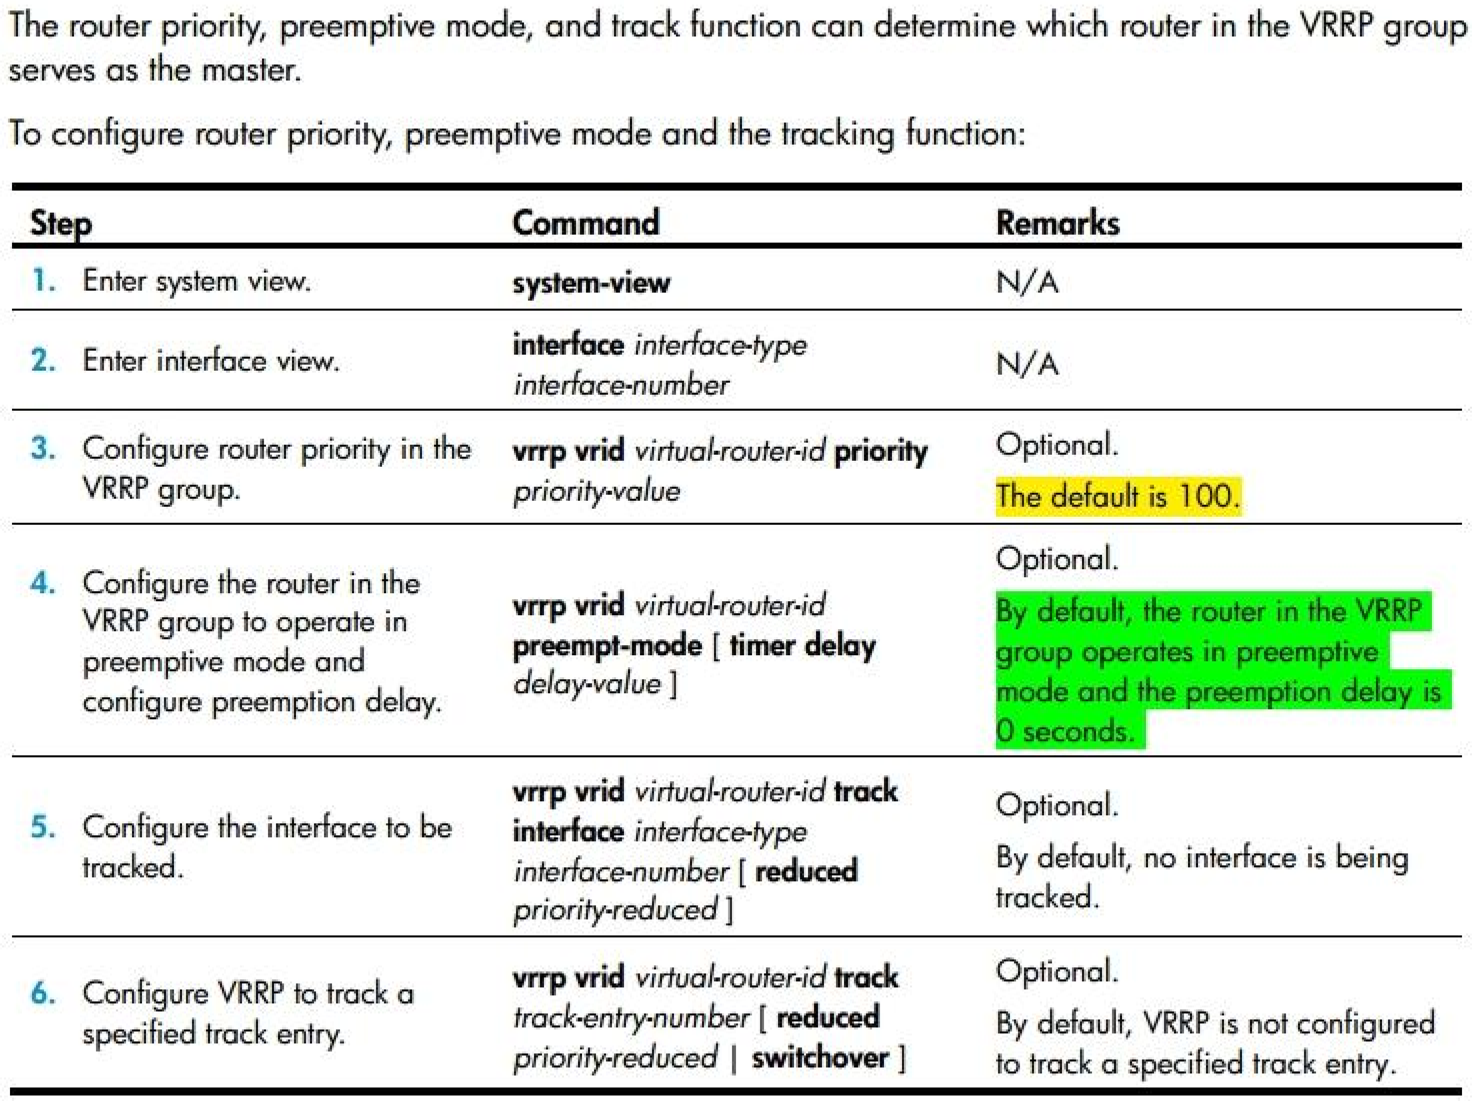

The router priority, preemptive mode, and track function can determine which router in the VRRP group serves as the master.

To configure router priority, preemptive mode and the tracking function:

Step	Command	Remarks
1. Enter system view.	**system-view**	N/A
2. Enter interface view.	**interface** *interface-type interface-number*	N/A
3. Configure router priority in the VRRP group.	**vrrp vrid** *virtual-router-id* **priority** *priority-value*	Optional. The default is 100.
4. Configure the router in the VRRP group to operate in preemptive mode and configure preemption delay.	**vrrp vrid** *virtual-router-id* **preempt-mode** [**timer delay** *delay-value*]	Optional. By default, the router in the VRRP group operates in preemptive mode and the preemption delay is 0 seconds.
5. Configure the interface to be tracked.	**vrrp vrid** *virtual-router-id* **track interface** *interface-type interface-number* [**reduced** *priority-reduced*]	Optional. By default, no interface is being tracked.
6. Configure VRRP to track a specified track entry.	**vrrp vrid** *virtual-router-id* **track** *track-entry-number* [**reduced** *priority-reduced* | **switchover**]	Optional. By default, VRRP is not configured to track a specified track entry.

QUESTION 21
View the exhibit. Based on the information provided in the exhibit, which IP address should this switch use as its router ID for OSPF?

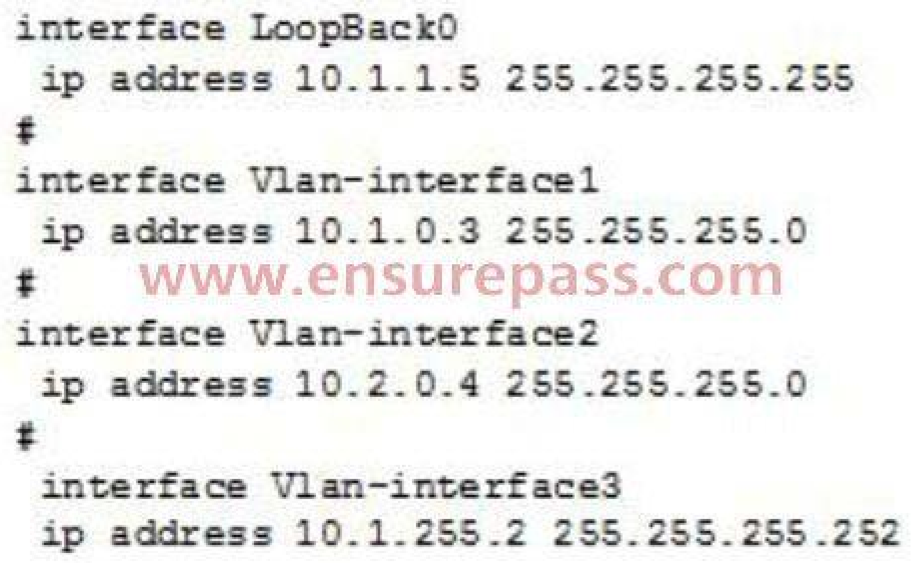

```
interface LoopBack0
 ip address 10.1.1.5 255.255.255.255
#
interface Vlan-interface1
 ip address 10.1.0.3 255.255.255.0
#
interface Vlan-interface2
 ip address 10.2.0.4 255.255.255.0
#
 interface Vlan-interface3
 ip address 10.1.255.2 255.255.255.252
```

A. 10.1.1.5
B. 10.2.0.4
C. 10.1.255.2
D. 10.1.0.3

Correct Answer: A
Section: OSPF
Explanation

Explanation/Reference:

QUESTION 22
What are functions of Type 9 LSAs in OSPFv3? (Select two.)

A. They enable the graceful restart of the OSPFv3 process.
B. They provide addressing information for links advertised in Type 1 and 2 LSAs.
C. They can trigger OSPF routing devices to run the shortest path first (SPF) algorithm again.
D. They advertise the prefixes for stub networks.
E. They help OSPF routing devices maintain their neighbors' state during the exchange of database information.

Correct Answer: BD
Section: OSPF
Explanation

Explanation/Reference:

QUESTION 23
A network administrator is configuring OSPF on an HP 5400 zl switch and has enabled OSPF globally. What is the minimum configuration required for the switch to begin sending OSPF hellos on a VLAN interface? (Select two.)

A. Assign the VLAN to an area.
B. Set the OSPF version.
C. Create an area globally.
D. Assign the switch an OSPF router ID.
E. Create a loopback interface.

Correct Answer: AC
Section: OSPF
Explanation

Explanation/Reference:

QUESTION 24
View the exhibit. Part of an OSPF routing table in an HP 5800 switch is shown in the exhibit. This switch is an internal switch in area 1. The network administrator views this routing table, and then reconfigures area 1 as a stub area on all relevant switches. No other changes to the OSPF domain are made. After this change takes effect, which routes shown in the exhibit will still exist in the OSPF routing table of this switch?

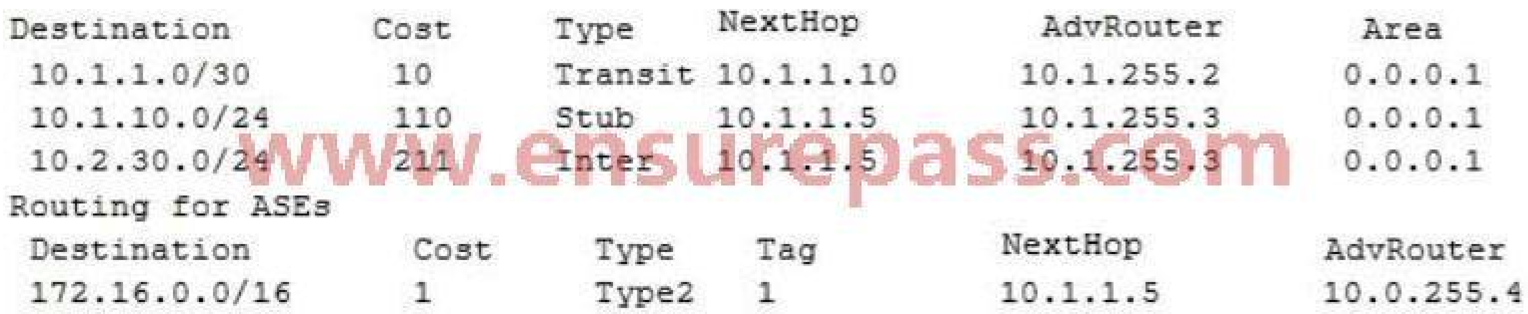

```
Destination        Cost     Type    NextHop          AdvRouter      Area
 10.1.1.0/30        10      Transit 10.1.1.10       10.1.255.2     0.0.0.1
 10.1.10.0/24       110     Stub    10.1.1.5        10.1.255.3     0.0.0.1
 10.2.30.0/24       211     Inter   10.1.1.5        10.1.255.3     0.0.0.1
Routing for ASEs
 Destination          Cost      Type    Tag         NextHop           AdvRouter
 172.16.0.0/16        1         Type2   1           10.1.1.5          10.0.255.4
```

A. All of the routes will still exist.

B. Only the routes to 10.1.1.0/30, 10.1.10.0/24, and 10.2.30.0/24 will still exist.
C. Only the routes to 10.1.1.0/30, 10.1.10.0/24, and 172.16.0.0/16 will still exist.
D. Only the routes to 10.1.1.0/30 and 10.1.10.0/24 will still exist.
E. Only the route to 10.1.10.0/24 will still exist.

Correct Answer: B
Section: OSPF
Explanation

Explanation/Reference:
OSPF allows certain areas to be configured as stub areas. External networks, such as those redistributed from other protocols into OSPF, are not allowed to be flooded into a stub area. Routing from these areas to the outside world is based on a default route. Configuring a stub area reduces the topological database size inside an area and reduces the memory requirements of routers inside that area.

QUESTION 25
View the exhibit. The HP 5406 zl switch configured with the VLANs shown in the exhibit receives a packet on a port that is an untagged member of VLAN 101. The packet is tagged with VLAN 20.
What will the switch do?

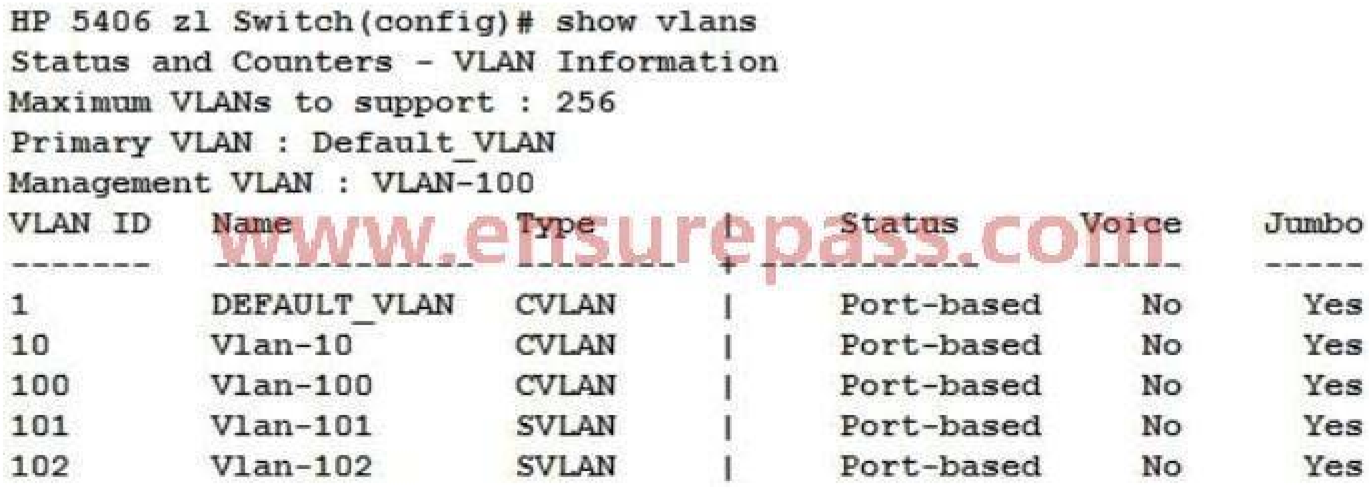

```
HP 5406 zl Switch(config)# show vlans
Status and Counters - VLAN Information
Maximum VLANs to support : 256
Primary VLAN : Default_VLAN
Management VLAN : VLAN-100
VLAN ID    Name          Type      |    Status      Voice    Jumbo
-------    ----------    --------  +  ----------    -----    -----
1          DEFAULT_VLAN  CVLAN     |    Port-based     No      Yes
10         Vlan-10       CVLAN     |    Port-based     No      Yes
100        Vlan-100      CVLAN     |    Port-based     No      Yes
101        Vlan-101      SVLAN     |    Port-based     No      Yes
102        Vlan-102      SVLAN     |    Port-based     No      Yes
```

A. The switch will insert another 802.1Q field and forward the packet.
B. The switch will drop the packet because the port must be a tagged member of VLAN 20.
C. The switch will drop the packet because the port must be a member of a C-VLAN.
D. The switch will remove the 802.1Q field and forward the packet.

Correct Answer: A
Section: VLANs
Explanation

Explanation/Reference:

How QinQ Works

Under QinQ, the provider network operates on a different VLAN space, independent of the VLANs that are used in the customer network as shown in Figure 8-2.

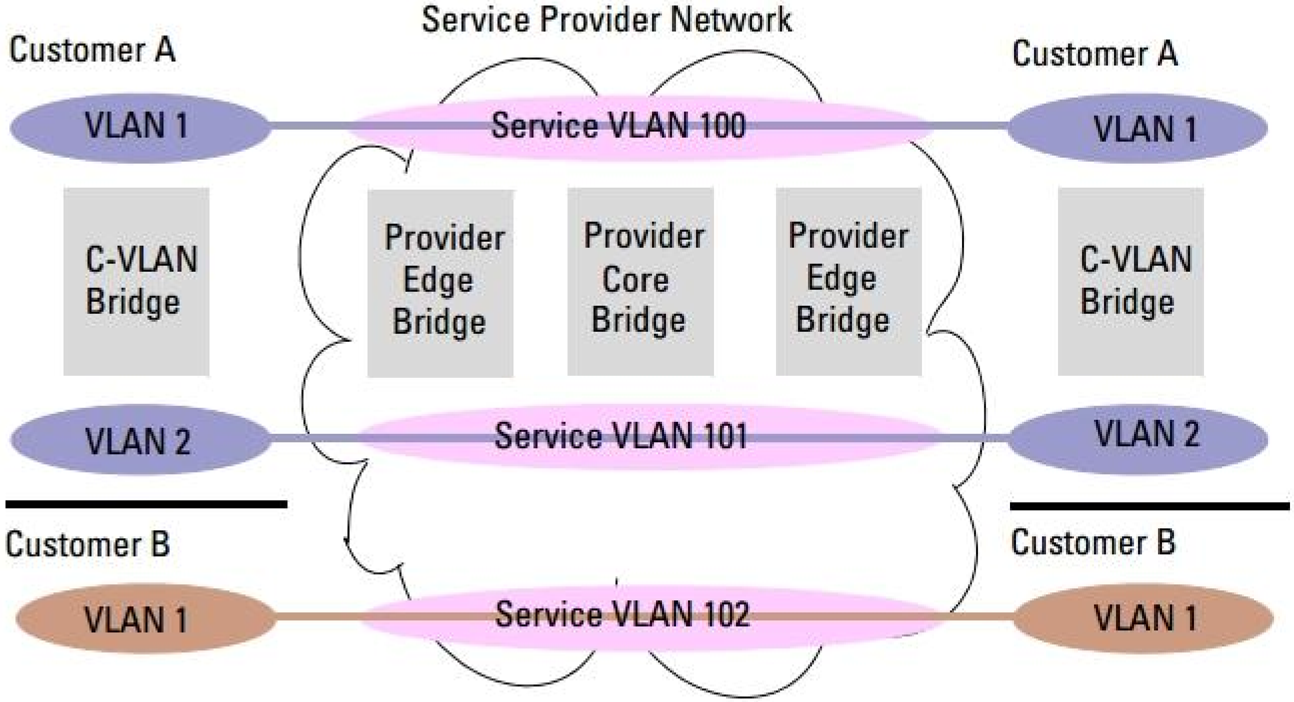

Figure 8-2. Example of VLANs in a QinQ Configuration

Customer VLANs (referred to as C-VLANs by the IEEE 802.1ad specification) are not used to make any forwarding decisions inside the provider network where customer frames get assigned to service VLANs (S-VLANs). Inside the provider cloud, frames are forwarded based on the S-VLAN tag only, while the C-VLAN tag remains shielded during data transmission. The S-VLAN tag is removed when the frame exits the provider network, restoring the original customer frame.

QUESTION 26
View the exhibits.

Exhibit 1
The frame has these characteristics:

VLAN tag: 2
Source MAC address: 000b-cdbb-3cc9
Source IP address: 10.1.4.12

Exhibit 2

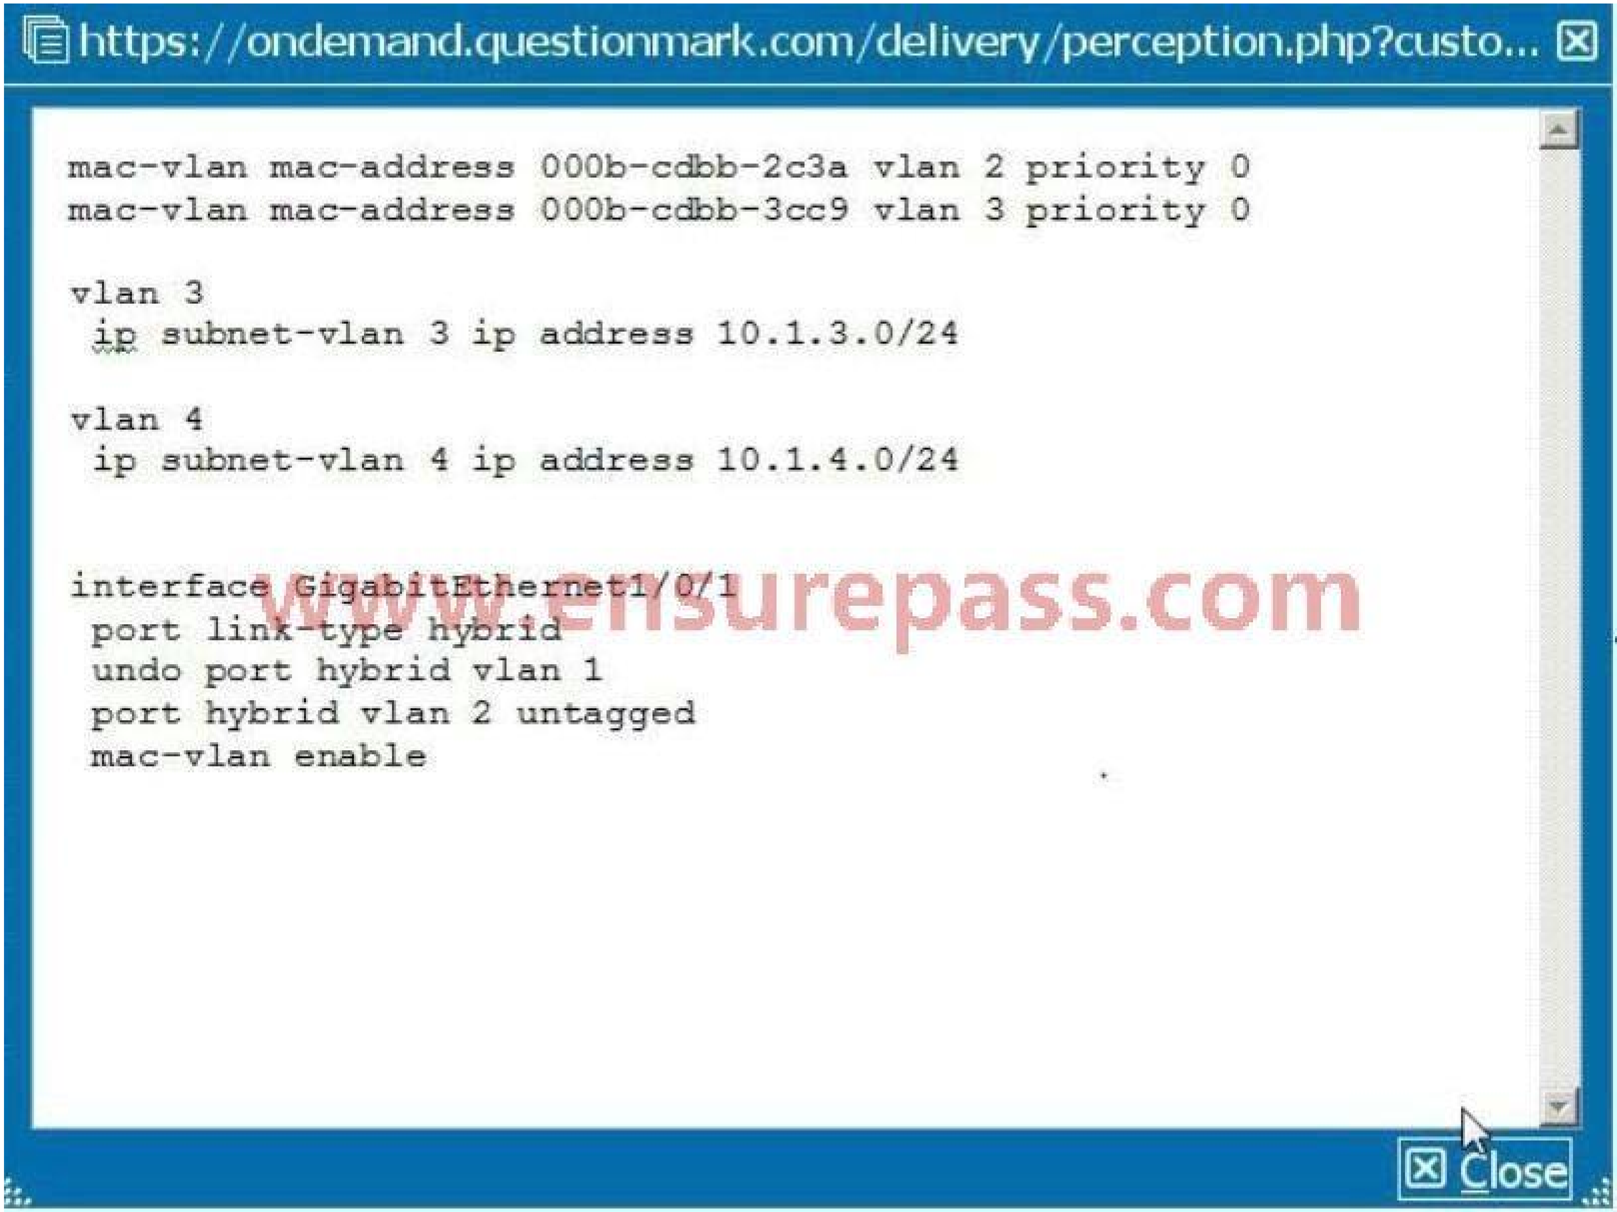
https://ondemand.questionmark.com/delivery/perception.php?custo...
Close

```
mac-vlan mac-address 000b-cdbb-2c3a vlan 2 priority 0
mac-vlan mac-address 000b-cdbb-3cc9 vlan 3 priority 0

vlan 3
 ip subnet-vlan 3 ip address 10.1.3.0/24

vlan 4
 ip subnet-vlan 4 ip address 10.1.4.0/24


interface GigabitEthernet1/0/1
 port link-type hybrid
 undo port hybrid vlan 1
 port hybrid vlan 2 untagged
 mac-vlan enable
```

The frame shown in Exhibit 1 arrives on port Gigabit Ethernet 1/0/1 on an HP 5800 switch. Based on the configuration shown in Exhibit 2, which statements are true? (Select two.)

A. The switch assigns the frame to VLAN 1.
B. The switch assigns the frame to VLAN 2.
C. The switch assigns the frame to VLAN 3.
D. The switch assigns the frame to VLAN 4.
E. The switch forwards the frame.
F. The switch drops the frame.

Correct Answer: CE
Section: VLANs
Explanation

Explanation/Reference:

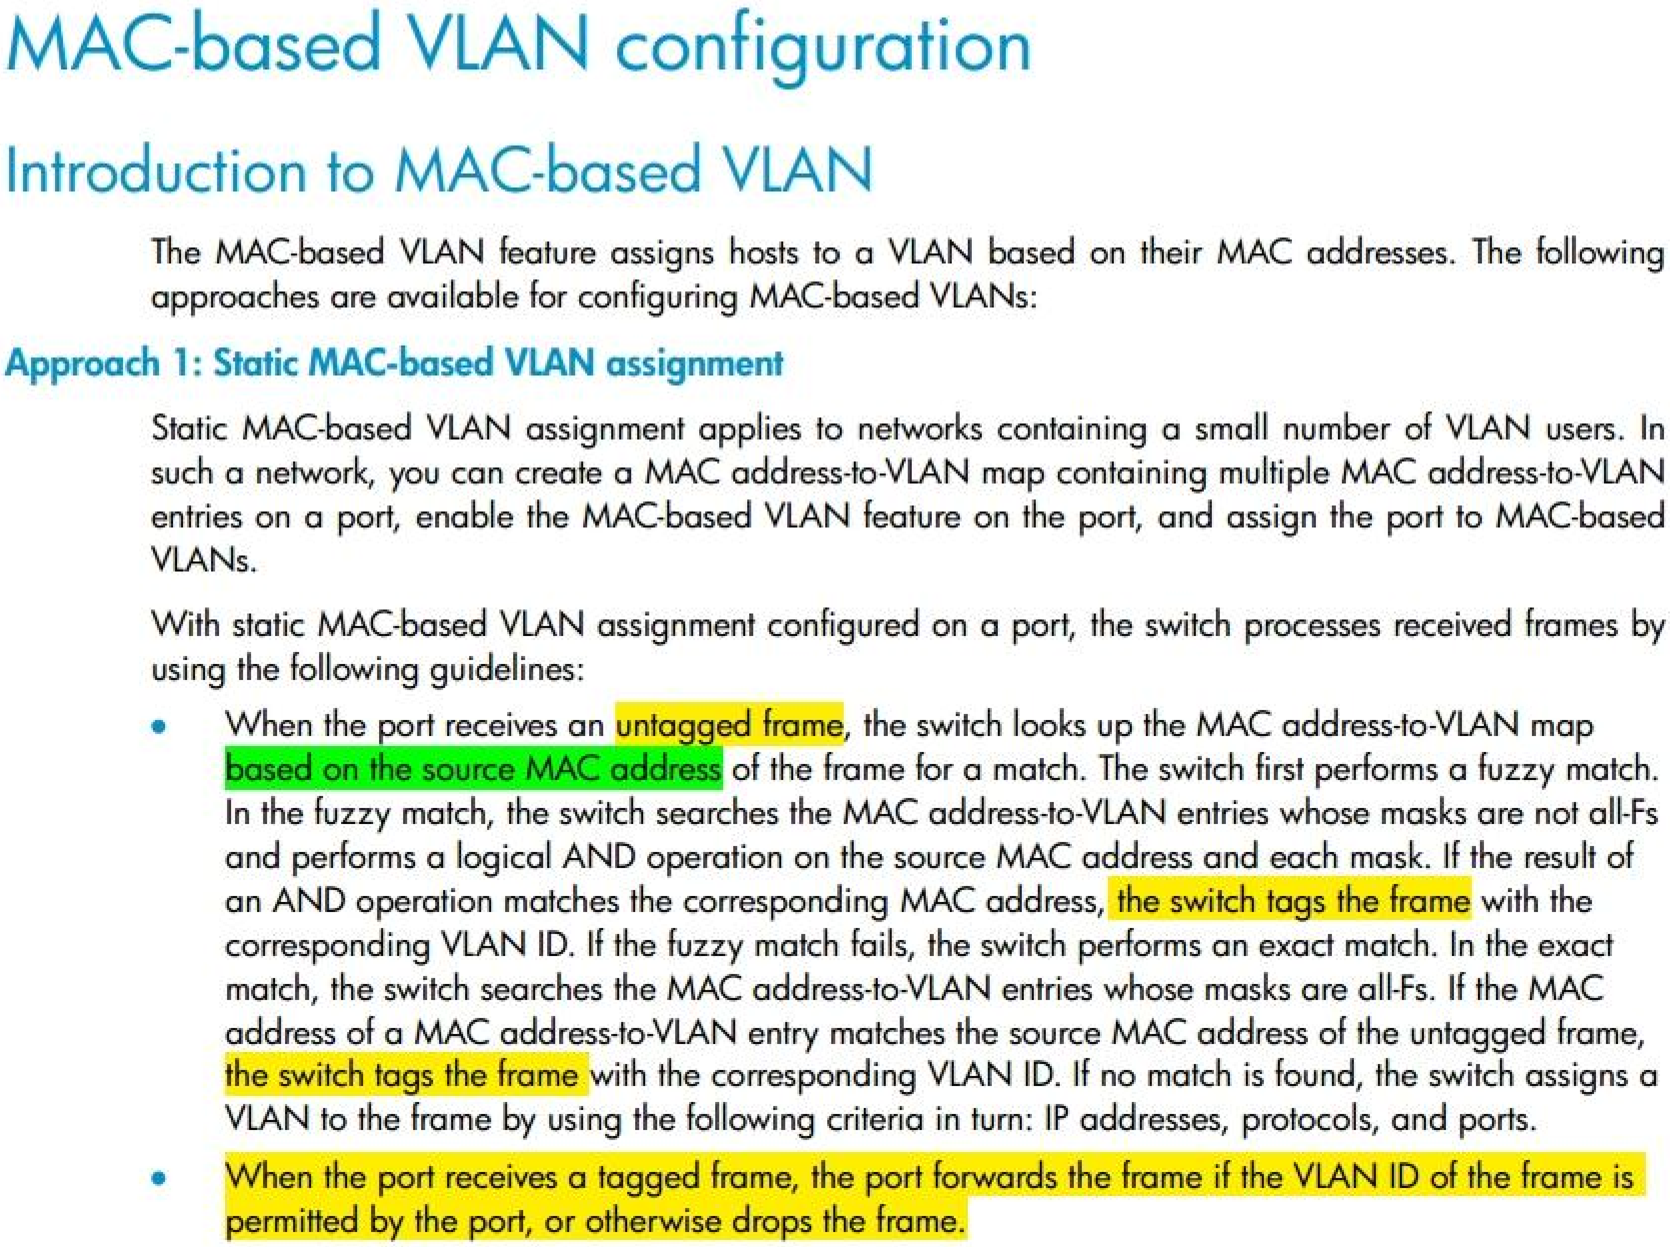

MAC-based VLAN configuration

Introduction to MAC-based VLAN

The MAC-based VLAN feature assigns hosts to a VLAN based on their MAC addresses. The following approaches are available for configuring MAC-based VLANs:

Approach 1: Static MAC-based VLAN assignment

Static MAC-based VLAN assignment applies to networks containing a small number of VLAN users. In such a network, you can create a MAC address-to-VLAN map containing multiple MAC address-to-VLAN entries on a port, enable the MAC-based VLAN feature on the port, and assign the port to MAC-based VLANs.

With static MAC-based VLAN assignment configured on a port, the switch processes received frames by using the following guidelines:

- When the port receives an untagged frame, the switch looks up the MAC address-to-VLAN map based on the source MAC address of the frame for a match. The switch first performs a fuzzy match. In the fuzzy match, the switch searches the MAC address-to-VLAN entries whose masks are not all-Fs and performs a logical AND operation on the source MAC address and each mask. If the result of an AND operation matches the corresponding MAC address, the switch tags the frame with the corresponding VLAN ID. If the fuzzy match fails, the switch performs an exact match. In the exact match, the switch searches the MAC address-to-VLAN entries whose masks are all-Fs. If the MAC address of a MAC address-to-VLAN entry matches the source MAC address of the untagged frame, the switch tags the frame with the corresponding VLAN ID. If no match is found, the switch assigns a VLAN to the frame by using the following criteria in turn: IP addresses, protocols, and ports.
- When the port receives a tagged frame, the port forwards the frame if the VLAN ID of the frame is permitted by the port, or otherwise drops the frame.

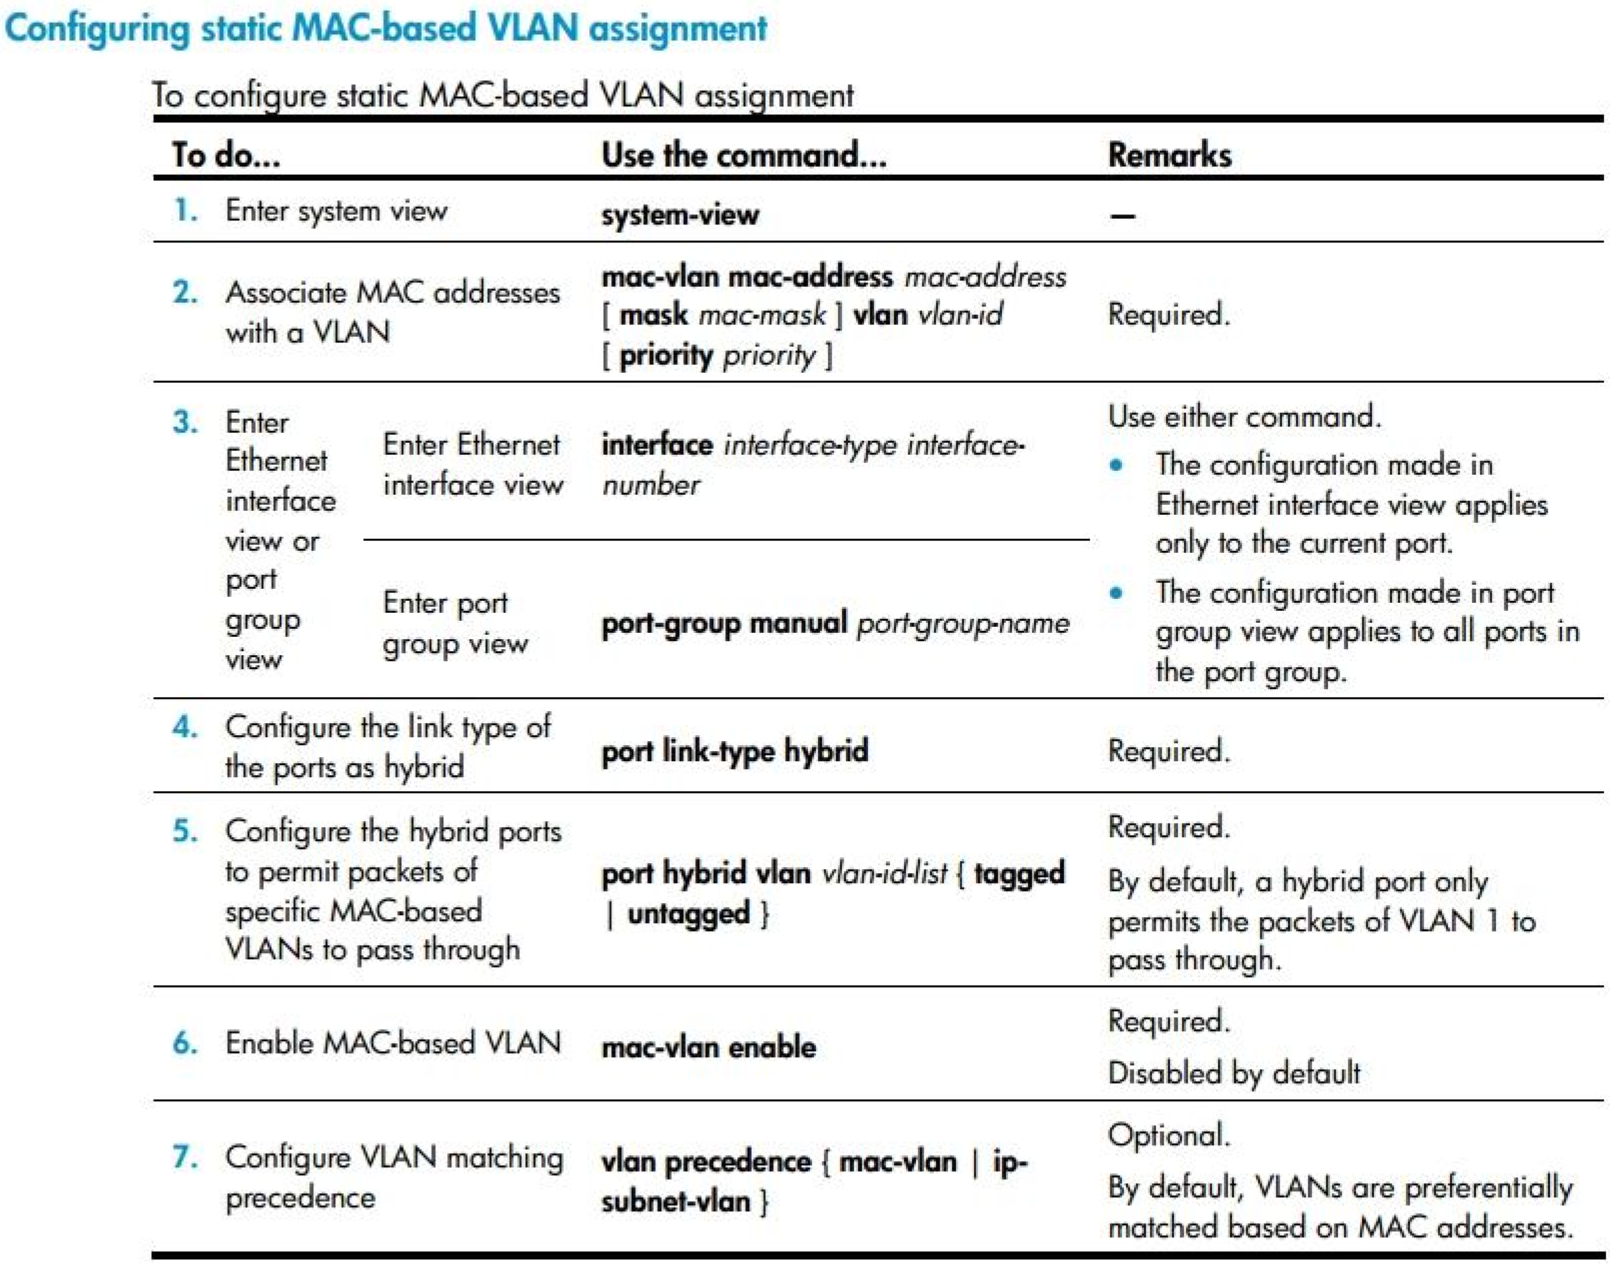

Configuring static MAC-based VLAN assignment

To configure static MAC-based VLAN assignment

To do...		Use the command...	Remarks
1. Enter system view		**system-view**	—
2. Associate MAC addresses with a VLAN		**mac-vlan mac-address** *mac-address* [**mask** *mac-mask*] **vlan** *vlan-id* [**priority** *priority*]	Required.
3. Enter Ethernet interface view or port group view	Enter Ethernet interface view	**interface** *interface-type interface-number*	Use either command. - The configuration made in Ethernet interface view applies only to the current port. - The configuration made in port group view applies to all ports in the port group.
	Enter port group view	**port-group manual** *port-group-name*	
4. Configure the link type of the ports as hybrid		**port link-type hybrid**	Required.
5. Configure the hybrid ports to permit packets of specific MAC-based VLANs to pass through		**port hybrid vlan** *vlan-id-list* { **tagged** | **untagged** }	Required. By default, a hybrid port only permits the packets of VLAN 1 to pass through.
6. Enable MAC-based VLAN		**mac-vlan enable**	Required. Disabled by default
7. Configure VLAN matching precedence		**vlan precedence** { **mac-vlan** | **ip-subnet-vlan** }	Optional. By default, VLANs are preferentially matched based on MAC addresses.

QUESTION 27
A service provider has configured an HP 5406 zl switch to support QinQ in mixed mode. Port A1 is a member of C-VLAN 100. Which usage is supported by this configuration?

A. Port A1 is an uplink port that forwards the customer's tunneled traffic to another switch in the service provider's network.
B. Port A1 is connected to a customer's switch, and the switch tunnels all VLAN traffic received on this port through the service provider's network.
C. Port A1 is connected to a device in the service provider's network and transmits traffic to and from that device.
D. Port A1 is connected to a customer's switch, and the switch tunnels C-VLAN 100 traffic received on this port through the service provider's network.

Correct Answer: B
Section: VLANs
Explanation

Explanation/Reference:
In this case, the configuration of switch 5406 is as follows:

```
(config)# qinq mixedvlan
(config)# svlan 100
(svlan-100)# untagged A1
(svlan-100)# tagged A2
(config)# int A1 qinq port-type customer-network
(config)# int A2 qinq port-type provider-network
```

Untagged port of SVLAN means that port A1 will receive traffic with 802 .1Q tag, i.e. the switch 5406 is connected to a customer's switch, and the switch tunnels all VLAN traffic received on this port through the service provider's network.

QUESTION 28
A company manages thousands of network infrastructure devices from several vendors with HP Intelligent Management Center (IMC). What advantage does the HP IMC distributed deployment model provide for this company?

A. The master server manages HP network infrastructure devices; slave servers manage non-HP devices.
B. The master server maintains the embedded database and Web browser interface; slave servers implement management functions.
C. A slave server duplicates the exact functions and database maintained on the master server and acts as a passive standby in case the master fails.
D. The slave servers offload some of the components from the master server to enhance performance.

Correct Answer: D
Section: IMC
Explanation

Explanation/Reference:

QUESTION 29
A company campus LAN requires high availability, so all edge switches have redundant links to the core. As part of the high availability design, all switches implement Multiple Spanning Tree Protocol (MSTP), and the core switches implement Virtual Router Redundancy Protocol (VRRP). Protocols are implemented to use links in the most efficient manner. What is the purpose of the link between the core switches during normal operation (when all links are up)?

A. The link acts as a backup link, blocked by MSTP in each instance but ready to transition to the forwarding state if necessary.
B. The link carries any traffic that arrives on the VRRP Backup to the VRRP Master.
C. The link does not carry any traffic but remains in forwarding state, ready to act as a backup in case another link fails.
D. The link carries VRRP messages between the VRRP Master and Backup on all VLANs.

Correct Answer: D
Section: High Availability
Explanation

Explanation/Reference:
Remember

QUESTION 30
View the exhibit. Topology 1 and topology 2 shown in the exhibit feature redundant gateways that implement Virtual Router Redundancy Protocol (VRRP). What is the advantage of topology 2 compared to topology 1?

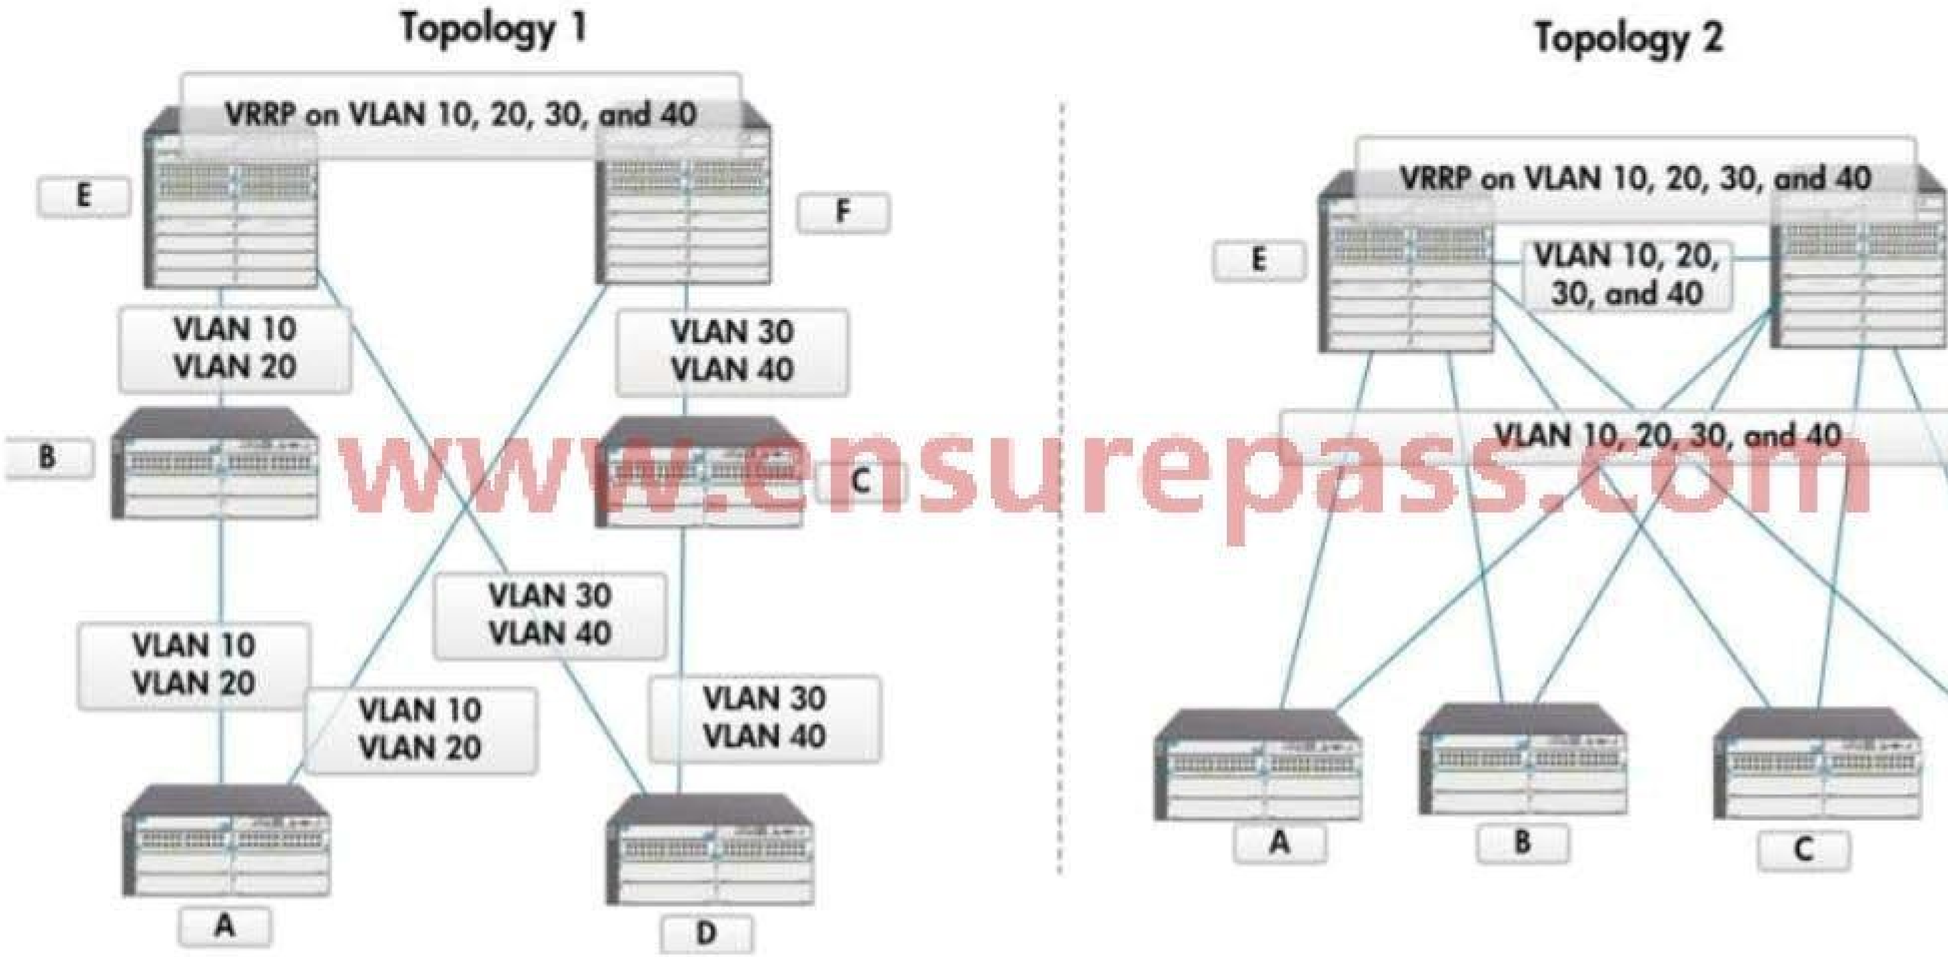

A. Topology 1 requires the implementation of other Layer 3 redundancy measures such as OSPF graceful restart.
B. Topology 2 makes it easier to load balance traffic between the two core switches.
C. Topology 2 provides true high availability if the link between the edge switches fails.
D. Topology 2 provides loop free logical topology and ensures maximum bandwidth usage.

Correct Answer: B
Section: High Availability
Explanation

Explanation/Reference:
Because of the edge switches traffic flowing directly into core switches, not flowing through all intermediate switches that causes them to overload, the traffic is easier to load balance between the two core switches.

QUESTION 31
A company is updating its campus LAN infrastructure, which includes a Unified Communications and Collaboration (UC&C) solution that features VoIP. What is the benefit of using HP 8206 zl core switches instead of HP 5406 zl switches?

A. The 8206 zl switches support LLDP-MED and the required TLVs.
B. The version 2 modules on the 8206 zl switches feature up to 8 configurable hardware queues.
C. The 8206 zl switches can provide advanced unicast and multicast routing without additional licensing.
D. The 8206 zl switches provide better high availability with non-stop switching and routing.

Correct Answer: D
Section: Other
Explanation

Explanation/Reference:
Remember

QUESTION 32
A campus LAN solution consists of HP Provision ASIC switches at the core and edge. Routing is implemented at the core. What is one common design principal for such a solution?

A. Create a different VLAN for each type of user or device on each different edge switch.
B. Avoid user-based VLAN assignments unless absolutely necessary.
C. Enable nonstop switching on edge switches to offload some processing from the core and improve performance.
D. Implement Multiple Spanning Tree Protocol (MSTP) on edge to core links for better utilization of redundant links' bandwidth.

Correct Answer: D
Section: Other
Explanation

Explanation/Reference:
Remember

QUESTION 33
View the exhibit. All switches in this figure are HP Provision ASIC switches. They are enabled to establish adjacency with neighbors on all of their VLAN interfaces. Besides these configurations, they are operating at the default OSPF settings. How can a network administrator ensure that routing switch B uses the 10G link to reach 10.1.10.0/24 if their direct links to routing switch A fails?

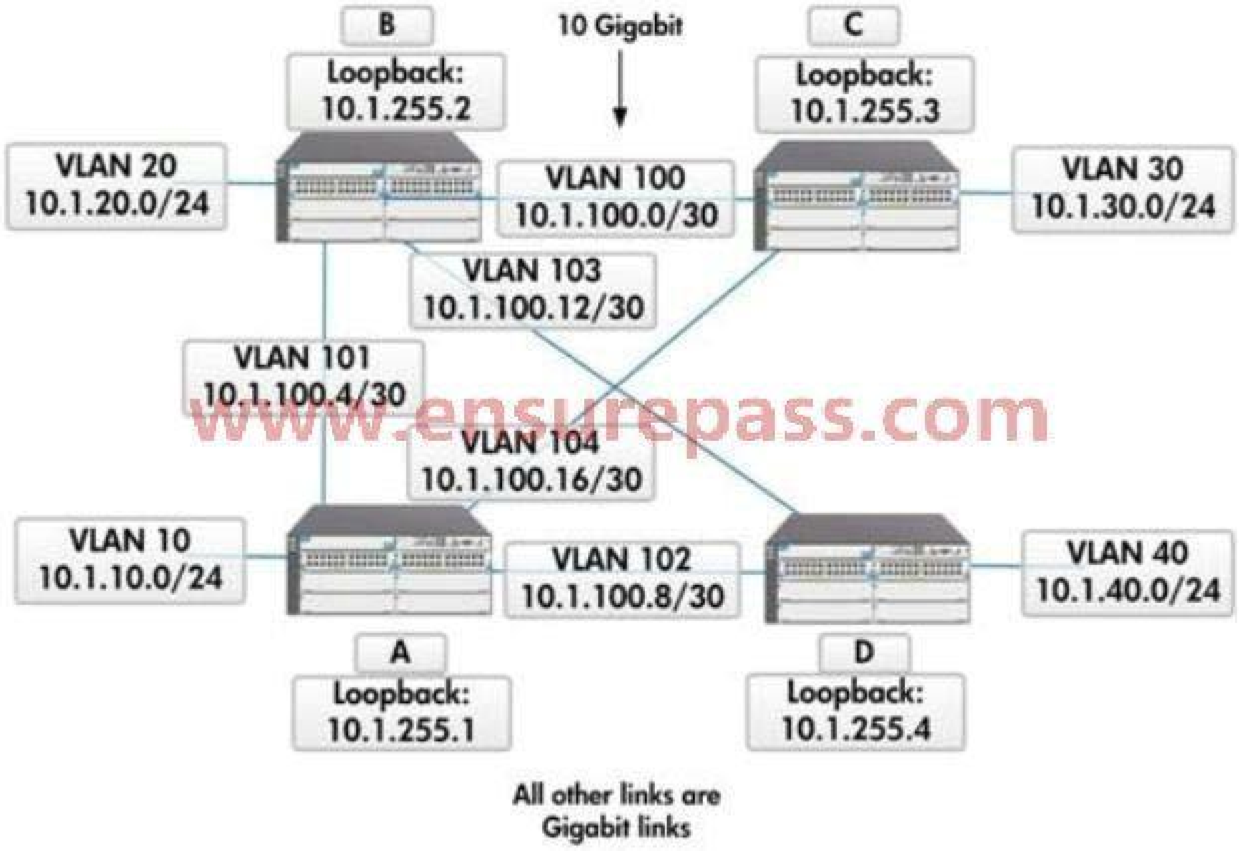

A. Set the VLAN 103 OSPF cost to 10.
B. Set the VLAN 100 bandwidth to 10000.
C. Set the VLAN 100 OSPF cost to 10.
D. Set the VLAN 103 bandwidth to 1000.
E. Set its reference bandwidth to 10000.

Correct Answer: E
Section: OSPF
Explanation

Explanation/Reference:
OSPF uses a reference bandwidth of 100 Mbps for cost calculation. The formula to calculate the cost is reference bandwidth divided by interface

bandwidth. For example, in the case of Ethernet, it is 100 Mbps / 10 Mbps = 10.

QUESTION 34
Before you enable dynamic IP lockdown on an 5400 zl switch, which feature must be enabled first?

A. connection-rate filtering
B. DHCP snooping
C. ARP protection
D. port security with eavesdropping protection

Correct Answer: B
Section: Other
Explanation

Explanation/Reference:

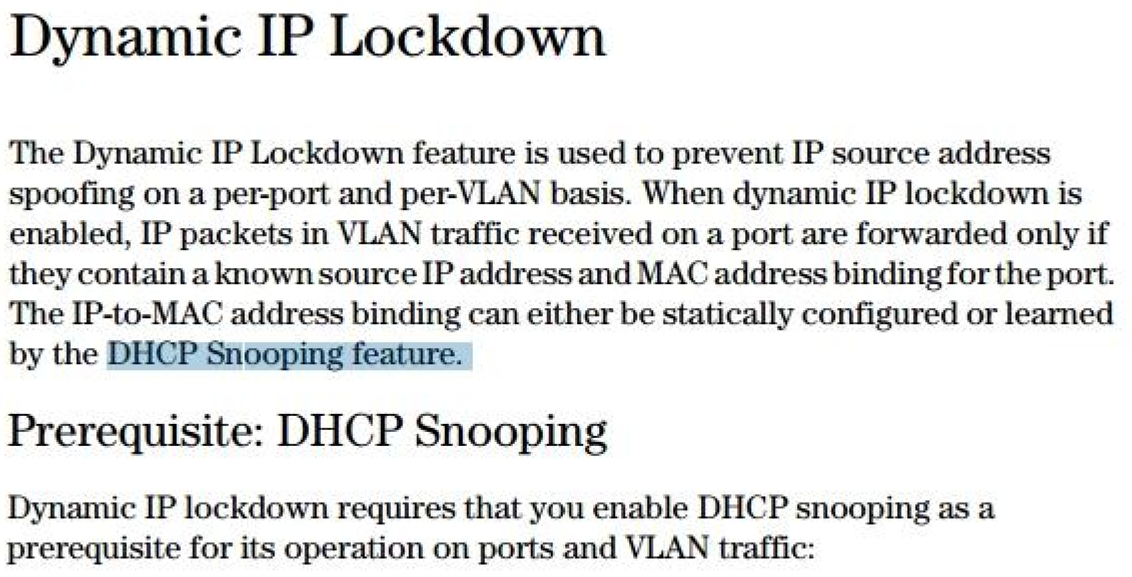

Dynamic IP Lockdown

The Dynamic IP Lockdown feature is used to prevent IP source address spoofing on a per-port and per-VLAN basis. When dynamic IP lockdown is enabled, IP packets in VLAN traffic received on a port are forwarded only if they contain a known source IP address and MAC address binding for the port. The IP-to-MAC address binding can either be statically configured or learned by the DHCP Snooping feature.

Prerequisite: DHCP Snooping

Dynamic IP lockdown requires that you enable DHCP snooping as a prerequisite for its operation on ports and VLAN traffic:

...

QUESTION 35
What is true about using EAP-Transport Layer Security (EAP-TLS) in the 802.1X process?

A. Supplicants use a non-reversible hash to submit all login credentials.
B. Supplicants are required to authenticate using digital certificates.

C. EAP-TLS provides a flexible security framework, which can be customized for each environment.

D. Supplicant and application servers are required to mutually authenticate with a username and password, which are encrypted before being transmitted.

Correct Answer: B
Section: Security
Explanation

Explanation/Reference:

QUESTION 36
A network administrator wants to configure the precise MAC addresses that are allowed to send traffic on a particular port on an HP 5400 zl Series switch. No other addresses can send traffic. Which task should the administrator complete?

A. Configure port security on the port in dynamic mode, specifying the correct MAC addresses.

B. Configure MAC lockout on the port, specifying the correct MAC addresses.

C. Configure port security on the port in configured mode, specifying the correct MAC addresses.

D. Configure MAC lockdown on the port, specifying the correct MAC addresses.

Correct Answer: D
Section: Security
Explanation

Explanation/Reference:

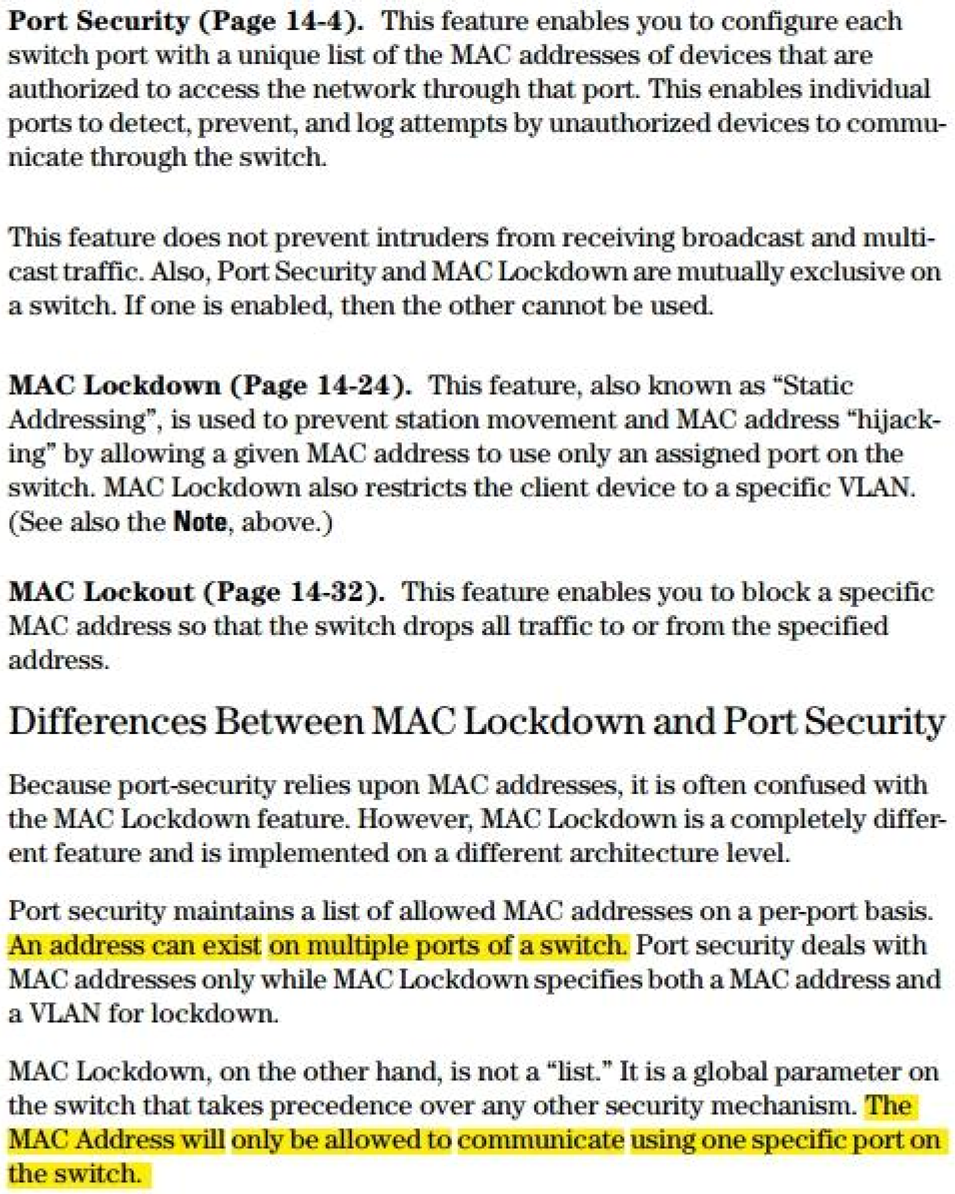

Port Security (Page 14-4). This feature enables you to configure each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch.

This feature does not prevent intruders from receiving broadcast and multicast traffic. Also, Port Security and MAC Lockdown are mutually exclusive on a switch. If one is enabled, then the other cannot be used.

MAC Lockdown (Page 14-24). This feature, also known as "Static Addressing", is used to prevent station movement and MAC address "hijacking" by allowing a given MAC address to use only an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. (See also the **Note**, above.)

MAC Lockout (Page 14-32). This feature enables you to block a specific MAC address so that the switch drops all traffic to or from the specified address.

Differences Between MAC Lockdown and Port Security

Because port-security relies upon MAC addresses, it is often confused with the MAC Lockdown feature. However, MAC Lockdown is a completely different feature and is implemented on a different architecture level.

Port security maintains a list of allowed MAC addresses on a per-port basis. An address can exist on multiple ports of a switch. Port security deals with MAC addresses only while MAC Lockdown specifies both a MAC address and a VLAN for lockdown.

MAC Lockdown, on the other hand, is not a "list." It is a global parameter on the switch that takes precedence over any other security mechanism. The MAC Address will only be allowed to communicate using one specific port on the switch.

QUESTION 37

Which statement accurately describes how network administrators can apply routed access control lists (RACLs) or VLAN access control lists (VACLs) on an HP 8200 zl switch?

A. A VACL applies to any traffic switched within a VLAN.
B. A RACL is applied at the global configuration level to filter all traffic routed on the switch.
C. A RACL controls all inbound or outbound routed traffic on a particular port.
D. A VACL applies to traffic routed within a particular VLAN.

Correct Answer: A
Section: Security
Explanation

Explanation/Reference:

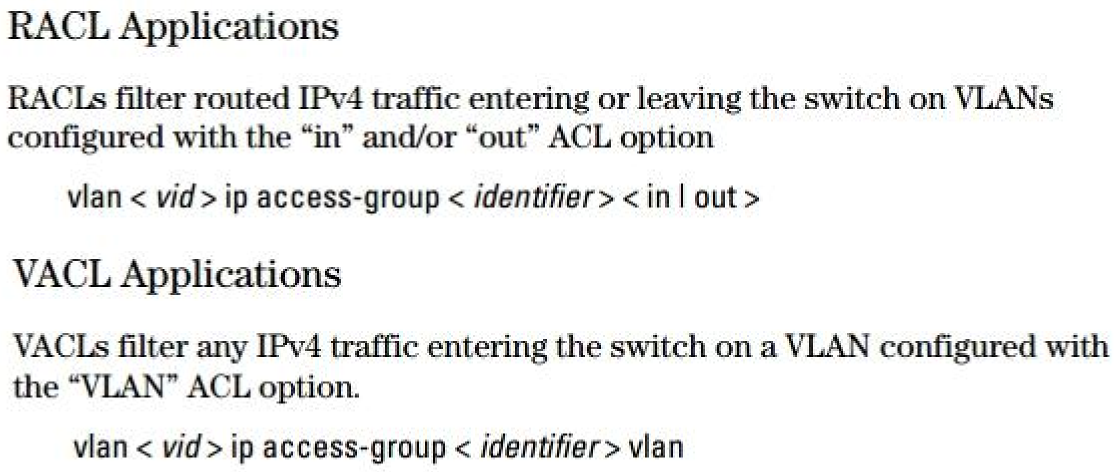

RACL Applications

RACLs filter routed IPv4 traffic entering or leaving the switch on VLANs configured with the "in" and/or "out" ACL option

```
vlan < vid > ip access-group < identifier > < in | out >
```

VACL Applications

VACLs filter any IPv4 traffic entering the switch on a VLAN configured with the "VLAN" ACL option.

```
vlan < vid > ip access-group < identifier > vlan
```

QUESTION 38
What is the first step a network administrator should complete to configure the first remote mirroring session between an HP 5400 zl and an HP 8200 zl switch?

A. Configure a GRE tunnel between the source and destination switch
B. Configure the remote mirror session on the destination switch
C. Add VLANs or ports to the remote mirror session
D. Configure the remote mirror session on the source switch

Correct Answer: B
Section: Other

Explanation

Explanation/Reference:

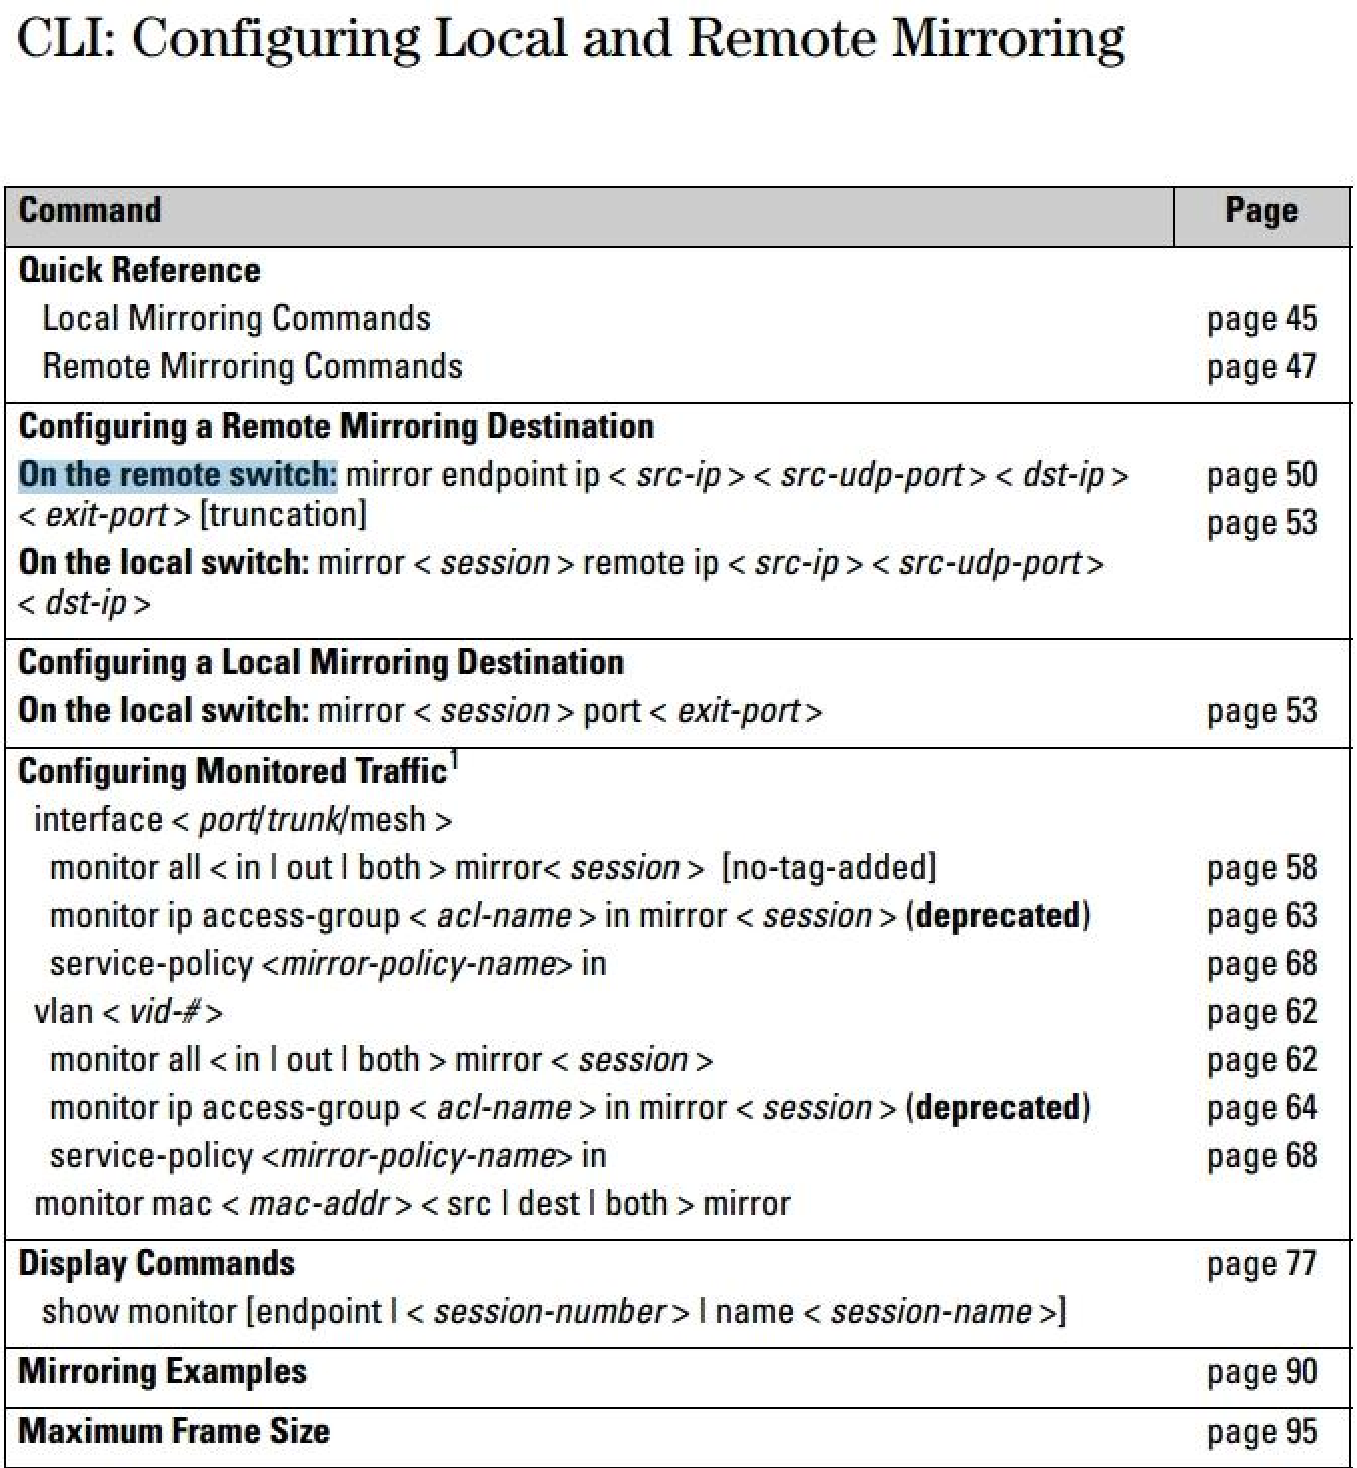

CLI: Configuring Local and Remote Mirroring

Command	Page
Quick Reference	
Local Mirroring Commands	page 45
Remote Mirroring Commands	page 47
Configuring a Remote Mirroring Destination	
On the remote switch: mirror endpoint ip < *src-ip* > < *src-udp-port* > < *dst-ip* > < *exit-port* > [truncation]	page 50
On the local switch: mirror < *session* > remote ip < *src-ip* > < *src-udp-port* > < *dst-ip* >	page 53
Configuring a Local Mirroring Destination	
On the local switch: mirror < *session* > port < *exit-port* >	page 53
Configuring Monitored Traffic[1]	
interface < *port*/*trunk*/mesh >	
monitor all < in | out | both > mirror< *session* > [no-tag-added]	page 58
monitor ip access-group < *acl-name* > in mirror < *session* > (**deprecated**)	page 63
service-policy <*mirror-policy-name*> in	page 68
vlan < *vid-#* >	page 62
monitor all < in | out | both > mirror < *session* >	page 62
monitor ip access-group < *acl-name* > in mirror < *session* > (**deprecated**)	page 64
service-policy <*mirror-policy-name*> in	page 68
monitor mac < *mac-addr* > < src | dest | both > mirror	
Display Commands	page 77
show monitor [endpoint | < *session-number* > | name < *session-name* >]	
Mirroring Examples	page 90
Maximum Frame Size	page 95

QUESTION 39
A network administrator enables OSPF on an HP Provision ASIC switch and creates area 0 and area 1. The network administrator then enters this command: Switch(ospf)# area 1 range 10.1.0.0/16. After other switches are configured and establish adjacencies, how will this command affect the final behavior of the switch?

A. The switch will become an area border router (ABR) capable of advertising summary LSAs.
B. The switch will send a single summary LSA into the backbone area for all networks in the 10.1.0.0/16 range.
C. The switch will send advertisements for networks within the 10.1.0.0/16 range into the backbone area.
D. The switch will send and accept OSPF messages on VLAN interfaces with the 10.1.0.0/16 range.

Correct Answer: B
Section: OSPF
Explanation

Explanation/Reference:

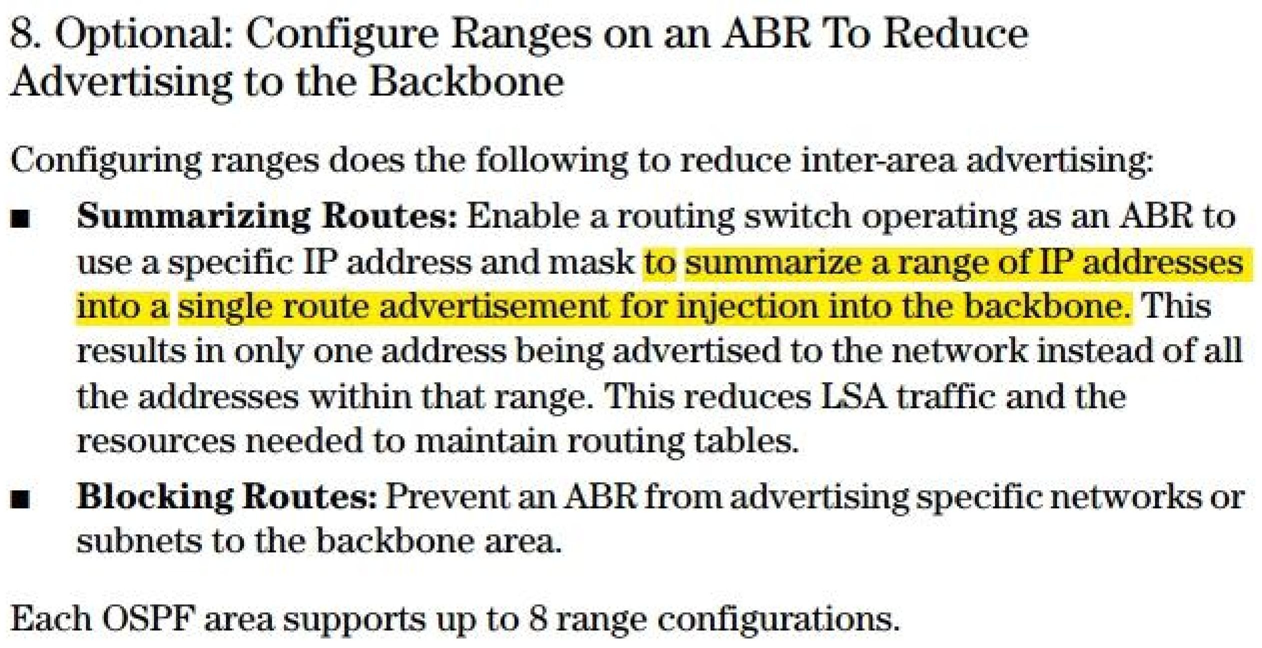

8. Optional: Configure Ranges on an ABR To Reduce Advertising to the Backbone

Configuring ranges does the following to reduce inter-area advertising:

- **Summarizing Routes:** Enable a routing switch operating as an ABR to use a specific IP address and mask to summarize a range of IP addresses into a single route advertisement for injection into the backbone. This results in only one address being advertised to the network instead of all the addresses within that range. This reduces LSA traffic and the resources needed to maintain routing tables.
- **Blocking Routes:** Prevent an ABR from advertising specific networks or subnets to the backbone area.

Each OSPF area supports up to 8 range configurations.

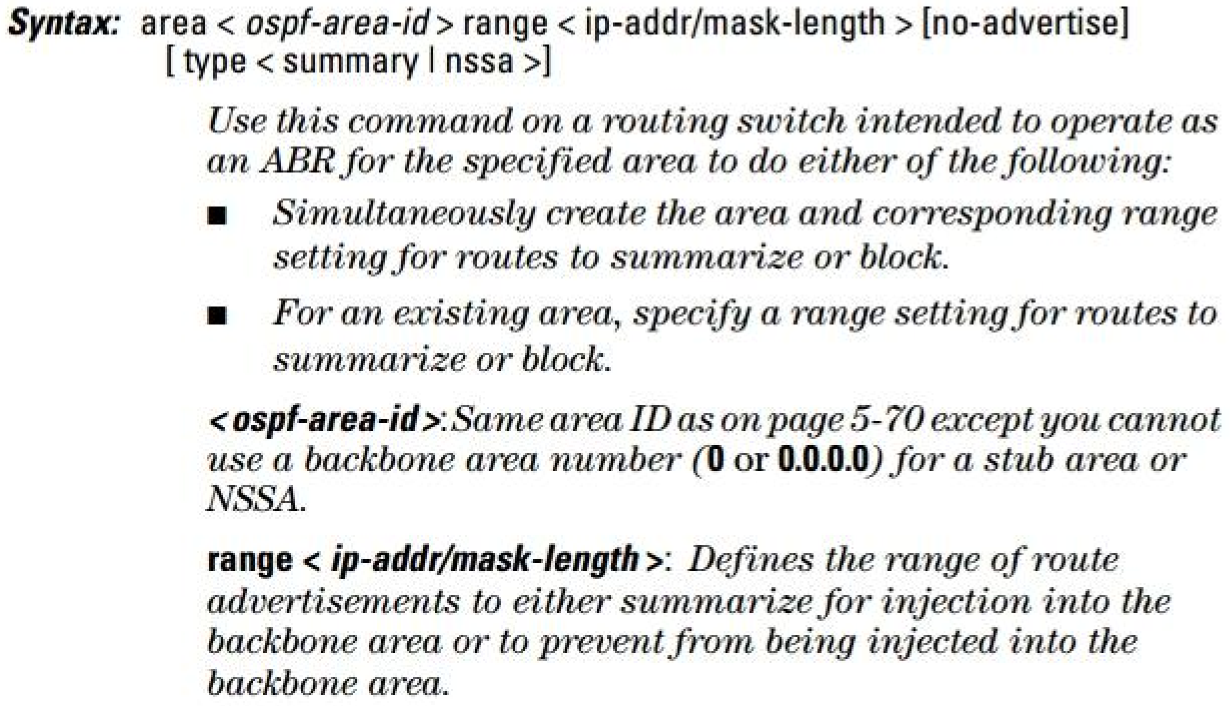

Syntax: area < *ospf-area-id* > range < ip-addr/mask-length > [no-advertise]
[type < summary | nssa >]

Use this command on a routing switch intended to operate as an ABR for the specified area to do either of the following:

- *Simultaneously create the area and corresponding range setting for routes to summarize or block.*
- *For an existing area, specify a range setting for routes to summarize or block.*

<*ospf-area-id*>:*Same area ID as on page 5-70 except you cannot use a backbone area number (***0** or **0.0.0.0***) for a stub area or NSSA.*

range < *ip-addr/mask-length* >: *Defines the range of route advertisements to either summarize for injection into the backbone area or to prevent from being injected into the backbone area.*

QUESTION 40
View the exhibit. Assume that this HP 10500 routing switch achieves adjacency with neighbors in Area 0 and Area 1. Based on the information provided in the exhibit, which statement accurately describes the behavior of this switch?

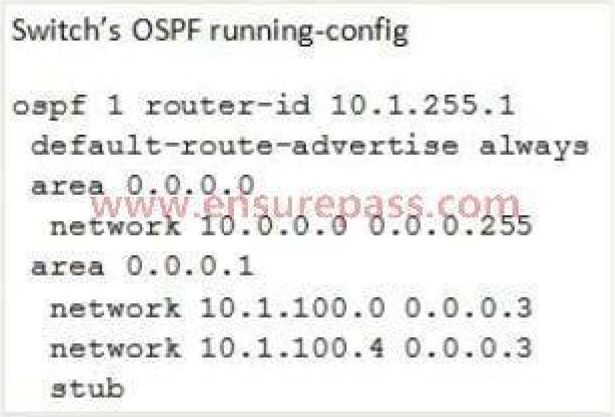

Switch's OSPF running-config

```
ospf 1 router-id 10.1.255.1
 default-route-advertise always
 area 0.0.0.0
  network 10.0.0.0 0.0.0.255
 area 0.0.0.1
  network 10.1.100.0 0.0.0.3
  network 10.1.100.4 0.0.0.3
  stub
```

A. It advertises a default route as an external route to OSPF neighbors in Area 0 and Area 1.
B. It advertises a default route as an external route to OSPF neighbors in Area 0 only.
C. It advertises a default route as an OSPF route to OSPF neighbors in Area 0 and Area 1.
D. It advertises a default route as an external route to OSPF neighbors in Area 0 and a default route as an OSPF route to OSPF neighbors in Area 1.
E. It advertises a default route as an OSPF route to OSPF neighbors in Area 1 only.

Correct Answer: A
Section: OSPF
Explanation

Explanation/Reference:

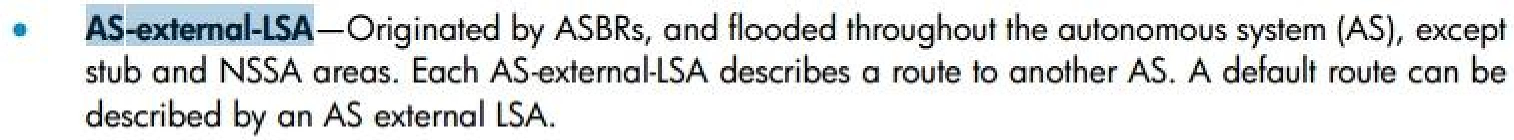

- **AS-external-LSA**—Originated by ASBRs, and flooded throughout the autonomous system (AS), except stub and NSSA areas. Each AS-external-LSA describes a route to another AS. A default route can be described by an AS external LSA.

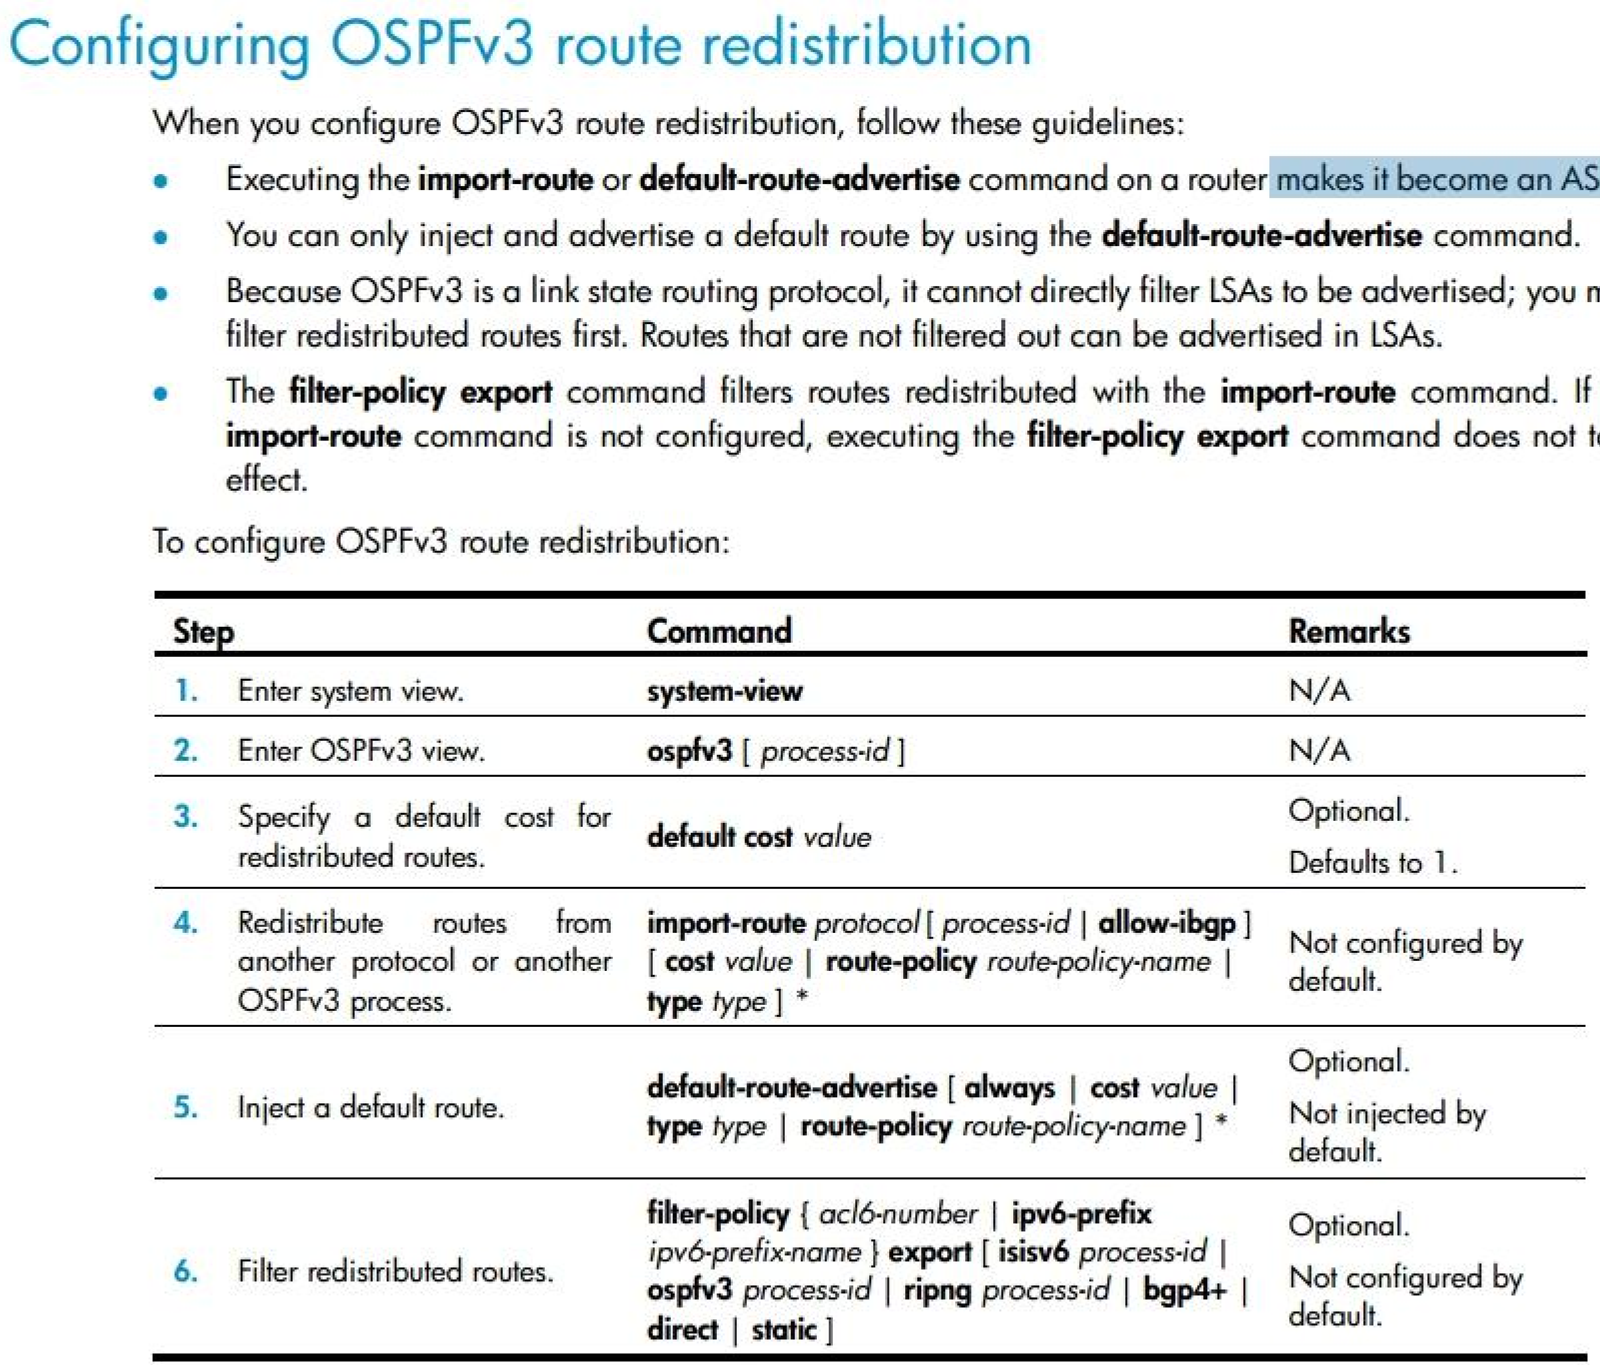

Configuring OSPFv3 route redistribution

When you configure OSPFv3 route redistribution, follow these guidelines:

- Executing the **import-route** or **default-route-advertise** command on a router makes it become an ASBR.
- You can only inject and advertise a default route by using the **default-route-advertise** command.
- Because OSPFv3 is a link state routing protocol, it cannot directly filter LSAs to be advertised; you must filter redistributed routes first. Routes that are not filtered out can be advertised in LSAs.
- The **filter-policy export** command filters routes redistributed with the **import-route** command. If the **import-route** command is not configured, executing the **filter-policy export** command does not take effect.

To configure OSPFv3 route redistribution:

Step	Command	Remarks
1. Enter system view.	**system-view**	N/A
2. Enter OSPFv3 view.	**ospfv3** [*process-id*]	N/A
3. Specify a default cost for redistributed routes.	**default cost** *value*	Optional. Defaults to 1.
4. Redistribute routes from another protocol or another OSPFv3 process.	**import-route** *protocol* [*process-id* | **allow-ibgp**] [**cost** *value* | **route-policy** *route-policy-name* | **type** *type*] *	Not configured by default.
5. Inject a default route.	**default-route-advertise** [**always** | **cost** *value* | **type** *type* | **route-policy** *route-policy-name*] *	Optional. Not injected by default.
6. Filter redistributed routes.	**filter-policy** { *acl6-number* | **ipv6-prefix** *ipv6-prefix-name* } **export** [**isisv6** *process-id* | **ospfv3** *process-id* | **ripng** *process-id* | **bgp4+** | **direct** | **static**]	Optional. Not configured by default.

QUESTION 41
View the exhibit. The OSPF routing table and OSPF settings on an HP 7500 Series routing switch are shown in the exhibit. Which networks does this switch advertise as summary LSAs in Area 0? (Select two.)

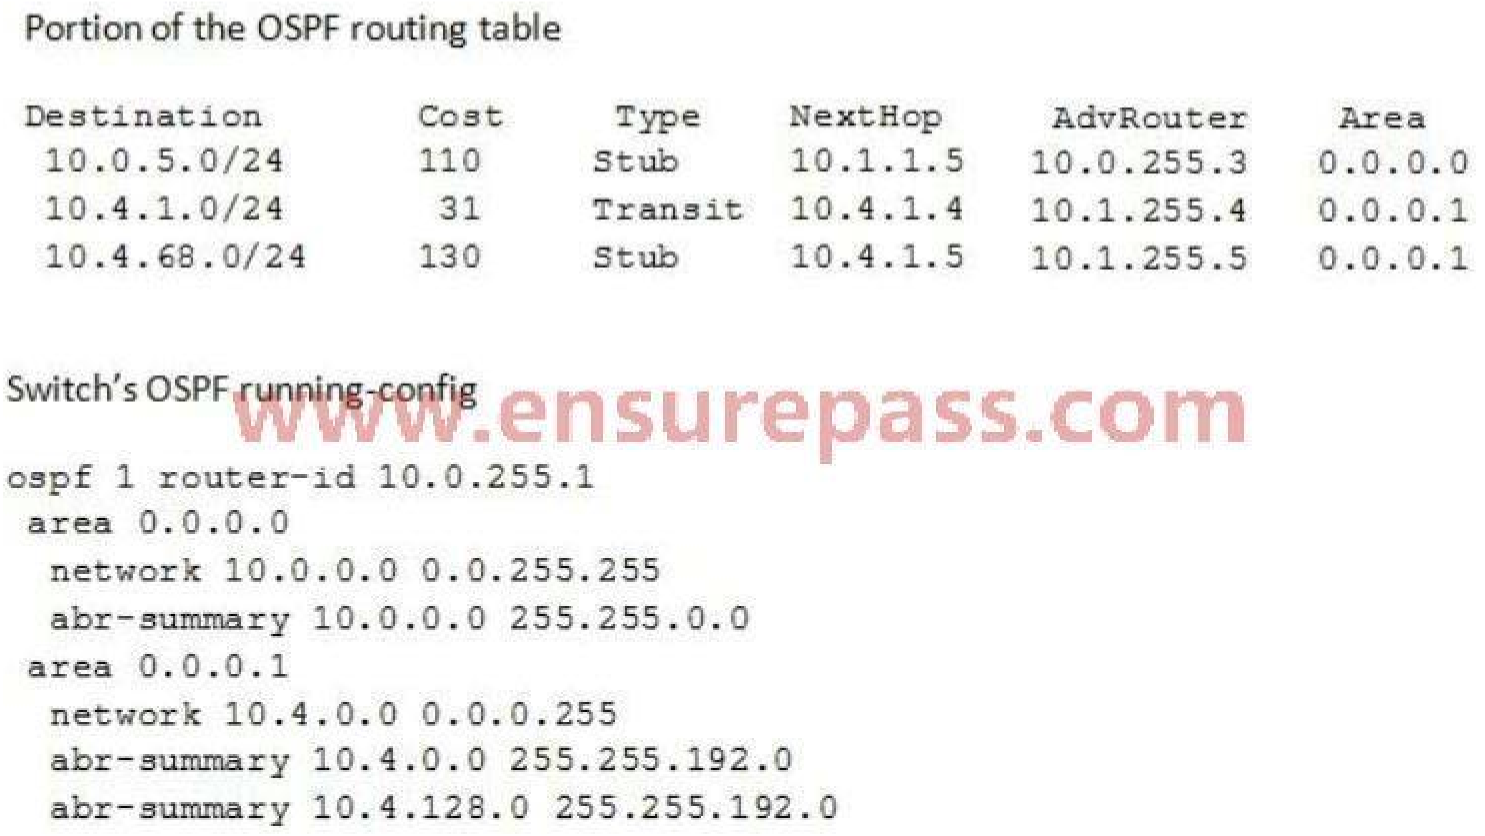

Portion of the OSPF routing table

```
Destination        Cost      Type      NextHop       AdvRouter      Area
 10.0.5.0/24        110      Stub      10.1.1.5      10.0.255.3     0.0.0.0
 10.4.1.0/24         31      Transit   10.4.1.4      10.1.255.4     0.0.0.1
 10.4.68.0/24       130      Stub      10.4.1.5      10.1.255.5     0.0.0.1
```

Switch's OSPF running-config

```
ospf 1 router-id 10.0.255.1
 area 0.0.0.0
  network 10.0.0.0 0.0.255.255
  abr-summary 10.0.0.0 255.255.0.0
 area 0.0.0.1
  network 10.4.0.0 0.0.0.255
  abr-summary 10.4.0.0 255.255.192.0
  abr-summary 10.4.128.0 255.255.192.0
```

A. 10.4.1.0/24
B. 10.0.5.0/24
C. 10.4.0.0/18
D. 10.4.128.0/18
E. 10.0.0.0/16
F. 10.4.68.0/24

Correct Answer: CD

Section: OSPF
Explanation

Explanation/Reference:

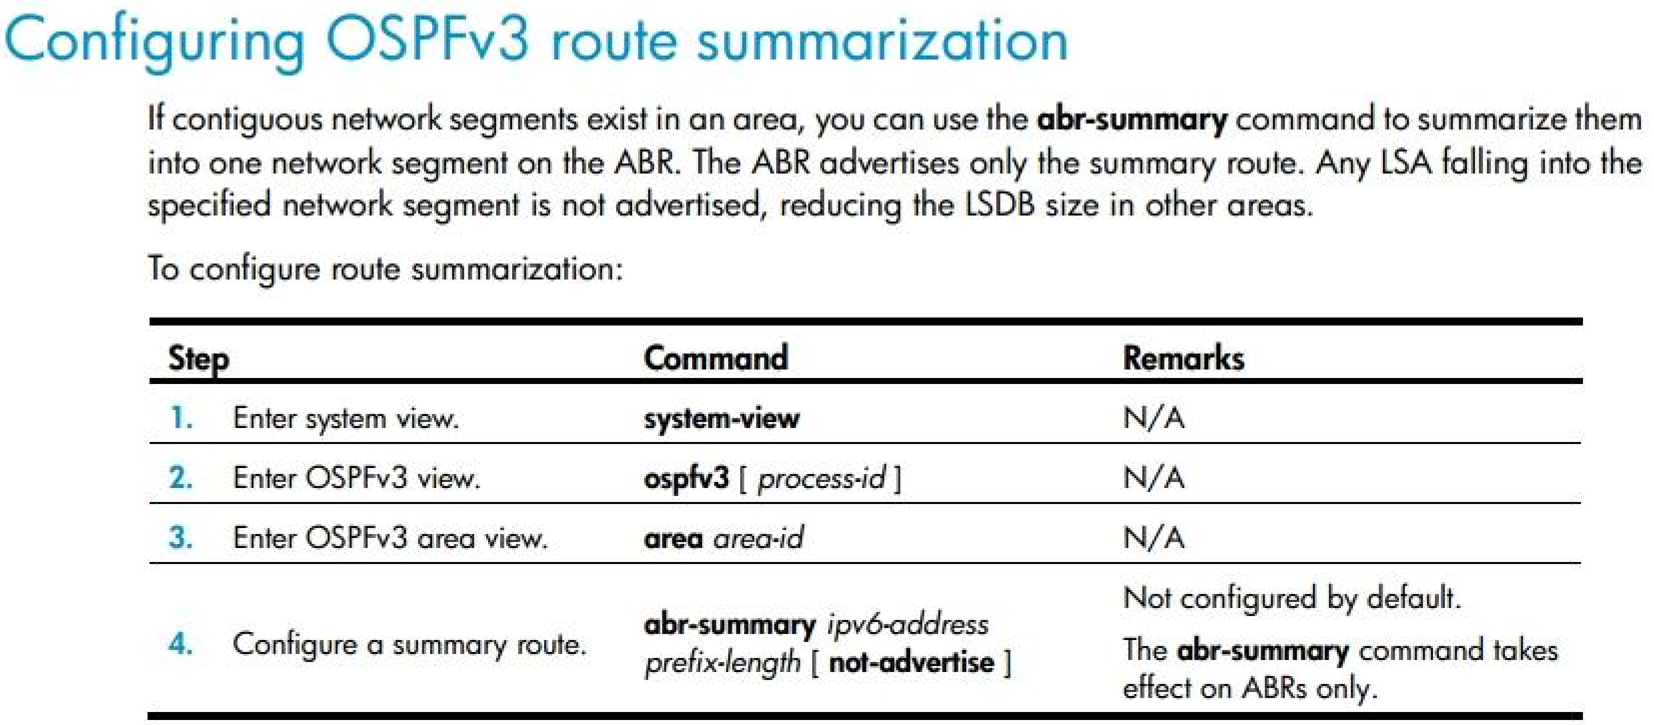

Configuring OSPFv3 route summarization

If contiguous network segments exist in an area, you can use the **abr-summary** command to summarize them into one network segment on the ABR. The ABR advertises only the summary route. Any LSA falling into the specified network segment is not advertised, reducing the LSDB size in other areas.

To configure route summarization:

Step		Command	Remarks
1.	Enter system view.	**system-view**	N/A
2.	Enter OSPFv3 view.	**ospfv3** [*process-id*]	N/A
3.	Enter OSPFv3 area view.	**area** *area-id*	N/A
4.	Configure a summary route.	**abr-summary** *ipv6-address prefix-length* [**not-advertise**]	Not configured by default. The **abr-summary** command takes effect on ABRs only.

QUESTION 42
View the exhibit.
Exhibit 1

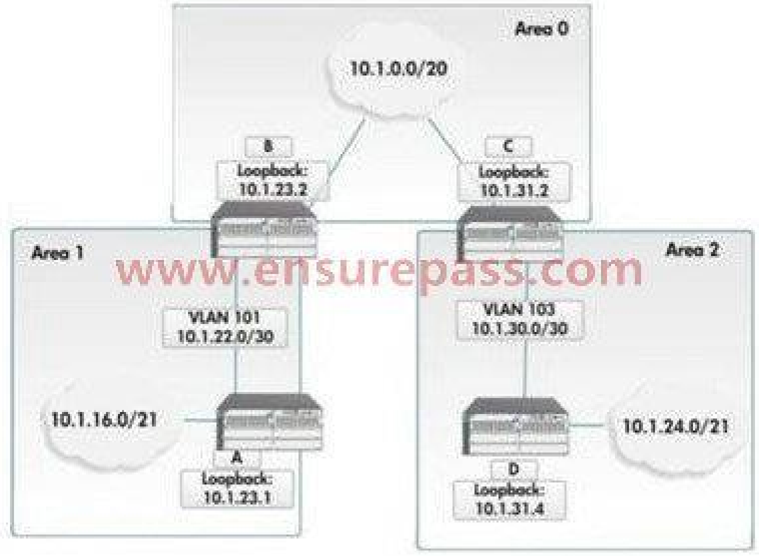

Exhibit 2

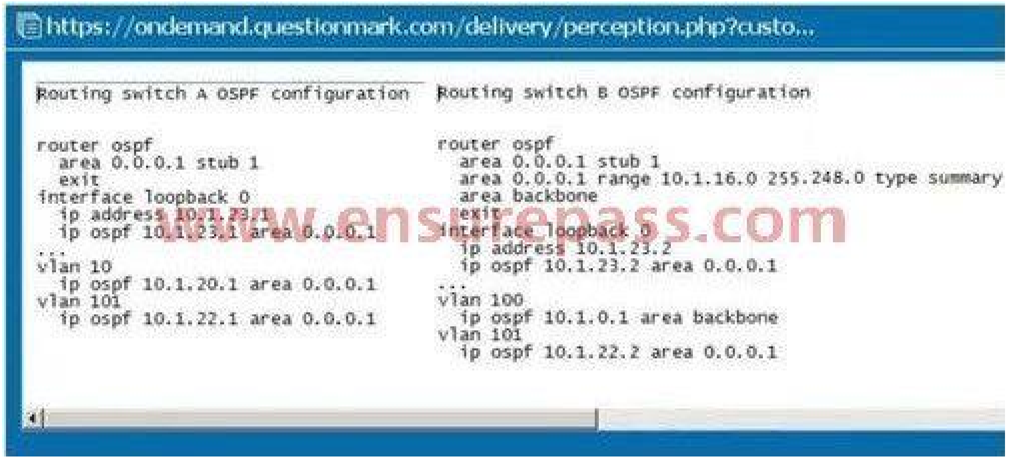

https://ondemand.questionmark.com/delivery/perception.php?custo...

```
Routing switch A OSPF configuration     Routing switch B OSPF configuration


router ospf                             router ospf
  area 0.0.0.1 stub 1                     area 0.0.0.1 stub 1
  exit                                    area 0.0.0.1 range 10.1.16.0 255.248.0 type summary
interface loopback 0                      area backbone
  ip address 10.1.23.1                    exit
  ip ospf 10.1.23.1 area 0.0.0.1        interface loopback 0
...                                       ip address 10.1.23.2
vlan 10                                   ip ospf 10.1.23.2 area 0.0.0.1
  ip ospf 10.1.20.1 area 0.0.0.1        ...
vlan 101                                vlan 100
  ip ospf 10.1.22.1 area 0.0.0.1          ip ospf 10.1.0.1 area backbone
                                        vlan 101
                                          ip ospf 10.1.22.2 area 0.0.0.1
```

Area 1 and Area 2 in the topology are associated with/21 subnets, shown in Exhibit 1. All of the switches in this topology are HP Provision ASIC switches, and they have established the appropriate adjacencies. In Area 1, one of the/24 subnets within the/21 range, 10.1.19.0/24, is not used. The network administrator must prevent unnecessary traffic from entering Area 1 if users attempt to reach an address in the unused range. What can the administrator do that meets this requirement without disrupting the solution?

A. Create a static route to 10.1.19.0/24 with the black hole option on routing switch A (and other area 1 routers).
B. Remove the stub definition for Area 1 on all routers and routing switches in Area 1.
C. Suppress the advertisement of the default route in Area 1 on routing switch B, the area border router (ABR).
D. Enter range 10.1.19.0/24 no-summary in the OSPF configuration context of routing switch B, the area border router (ABR).

Correct Answer: A
Section: OSPF
Explanation

Explanation/Reference:

QUESTION 43
View the exhibit. In the example network shown in the exhibit, switch A loses power and becomes unavailable. Which switch ports have the designated role in instance 1 when the MSTP network converges again?

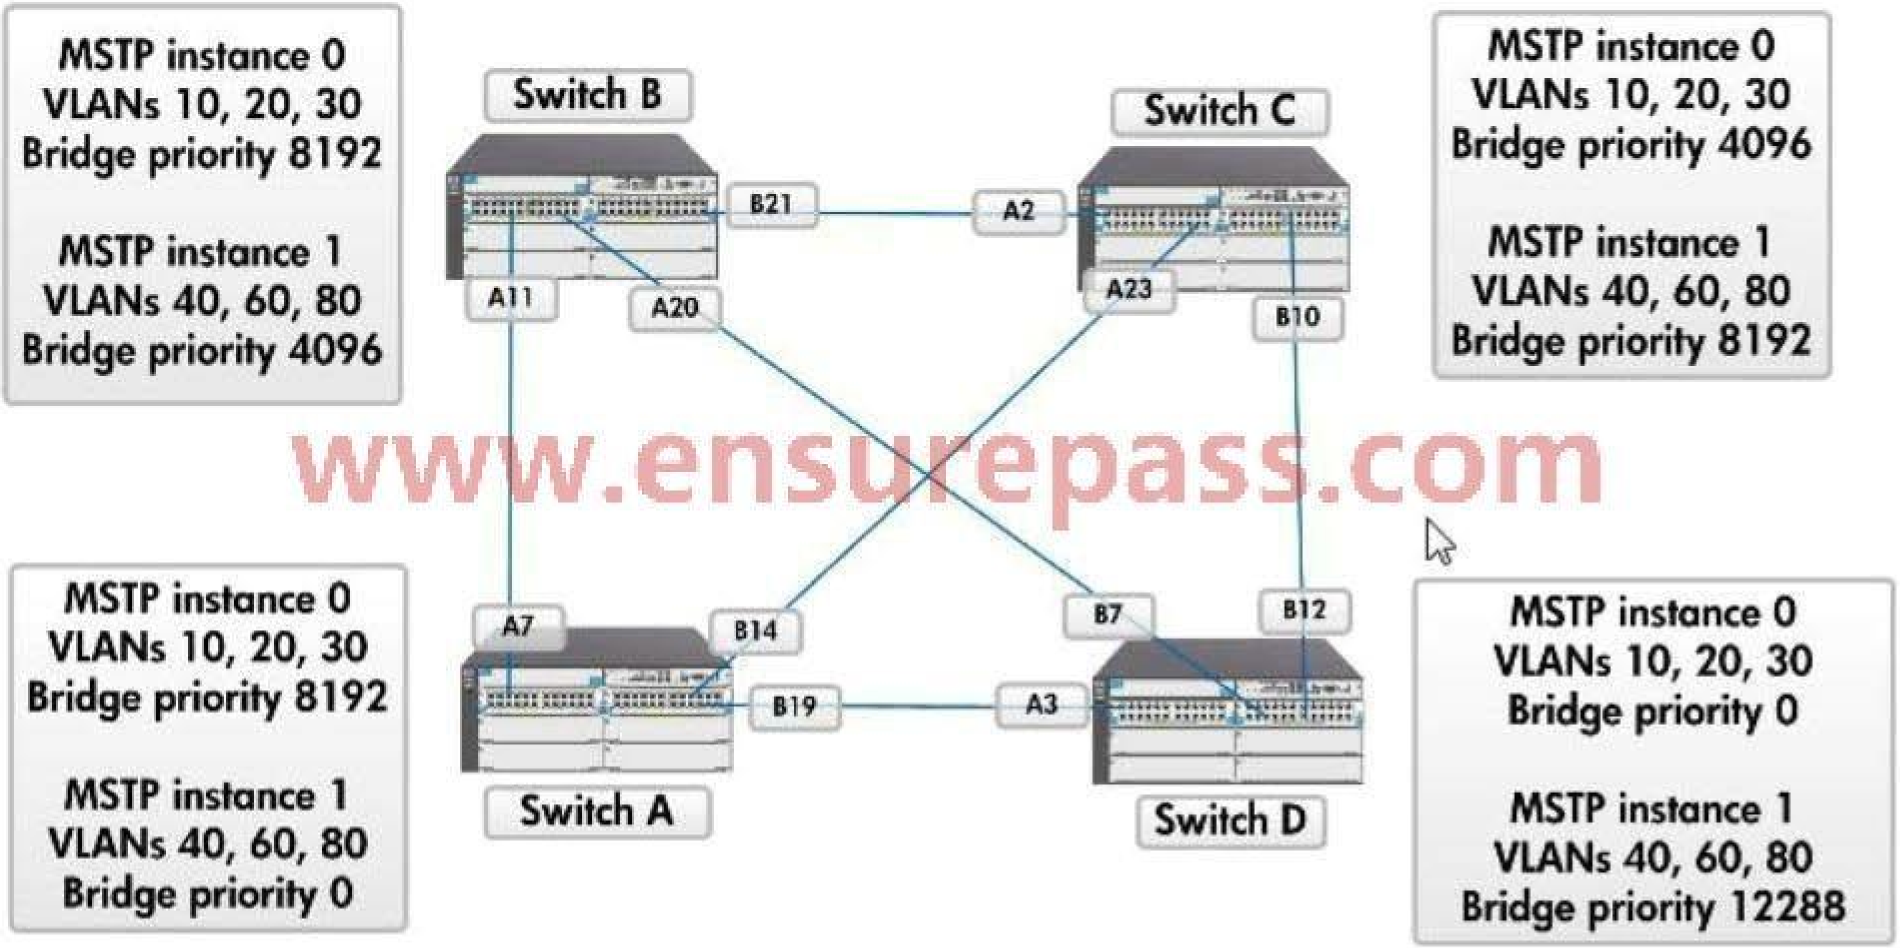

A. Ports A2 and B10 on switch C
B. Ports B7 on switch D and B10 on switch C
C. Ports A20 and B21 on switch B
D. Ports B7 and B12 on switch D

Correct Answer: C
Section: MSTP
Explanation

Explanation/Reference:

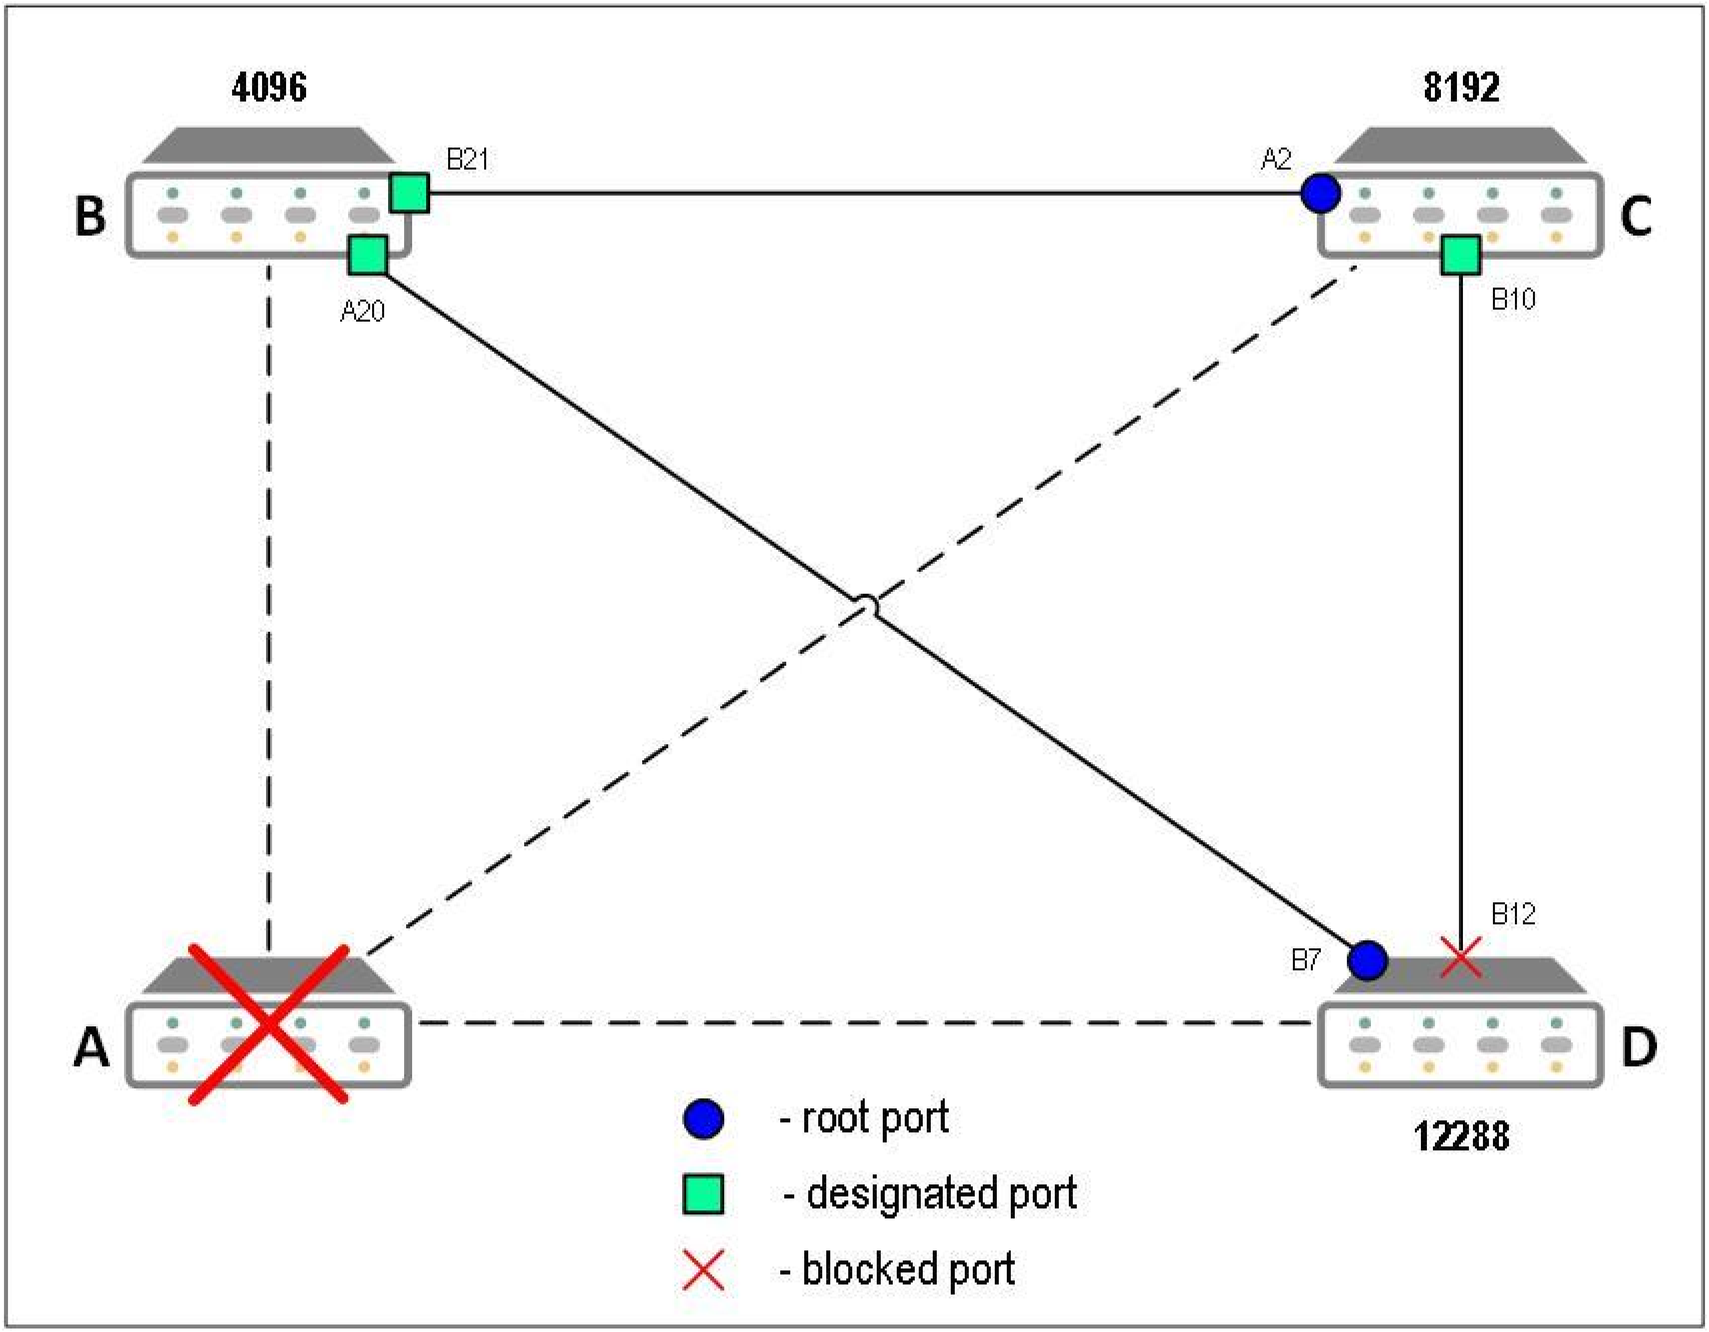
4096
8192
B
B21
A2
C
A20
B10
B12
B7
A
D
12288
● - root port
🟩 - designated port
✕ - blocked port

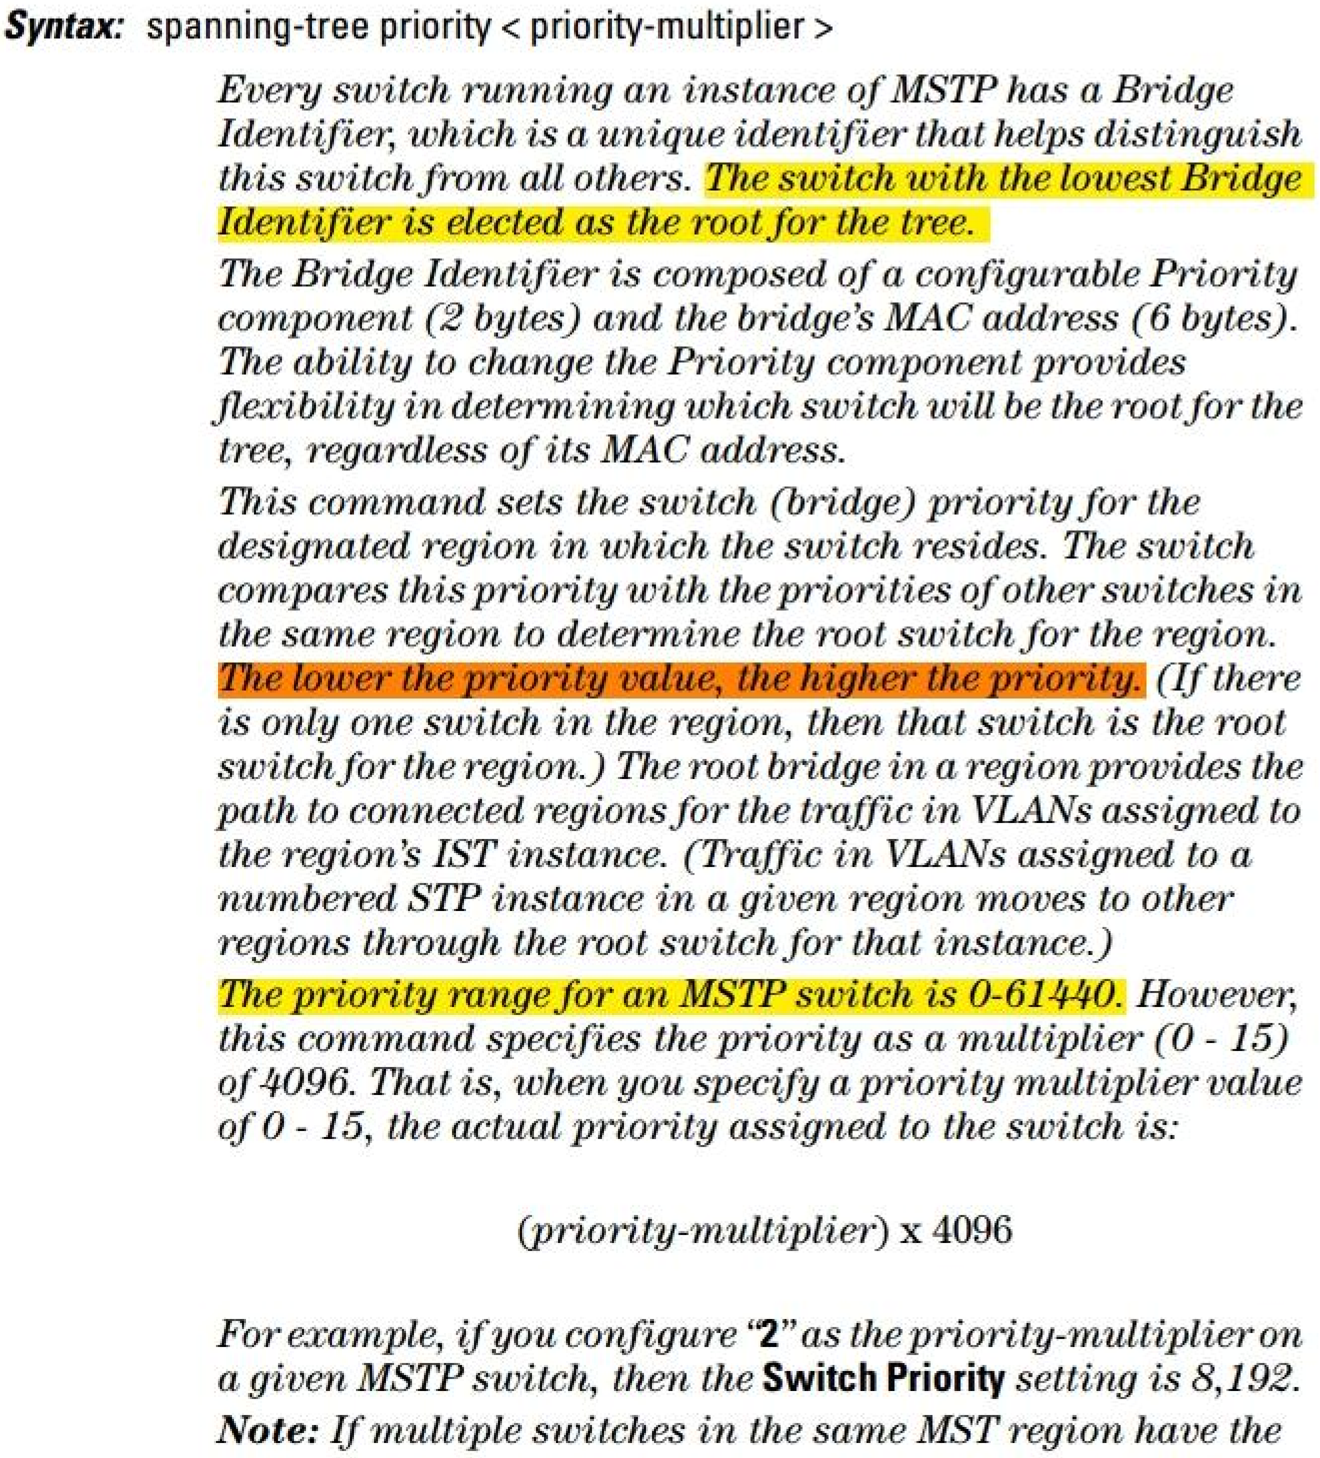

Syntax: spanning-tree priority < priority-multiplier >

Every switch running an instance of MSTP has a Bridge Identifier, which is a unique identifier that helps distinguish this switch from all others. The switch with the lowest Bridge Identifier is elected as the root for the tree.

The Bridge Identifier is composed of a configurable Priority component (2 bytes) and the bridge's MAC address (6 bytes). The ability to change the Priority component provides flexibility in determining which switch will be the root for the tree, regardless of its MAC address.

This command sets the switch (bridge) priority for the designated region in which the switch resides. The switch compares this priority with the priorities of other switches in the same region to determine the root switch for the region. The lower the priority value, the higher the priority. (If there is only one switch in the region, then that switch is the root switch for the region.) The root bridge in a region provides the path to connected regions for the traffic in VLANs assigned to the region's IST instance. (Traffic in VLANs assigned to a numbered STP instance in a given region moves to other regions through the root switch for that instance.)

The priority range for an MSTP switch is 0-61440. However, this command specifies the priority as a multiplier (0 - 15) of 4096. That is, when you specify a priority multiplier value of 0 - 15, the actual priority assigned to the switch is:

(priority-multiplier) x 4096

*For example, if you configure "***2***" as the priority-multiplier on a given MSTP switch, then the* **Switch Priority** *setting is 8,192.*

Note: *If multiple switches in the same MST region have the*

QUESTION 44
View the exhibit.

Exhibit 1

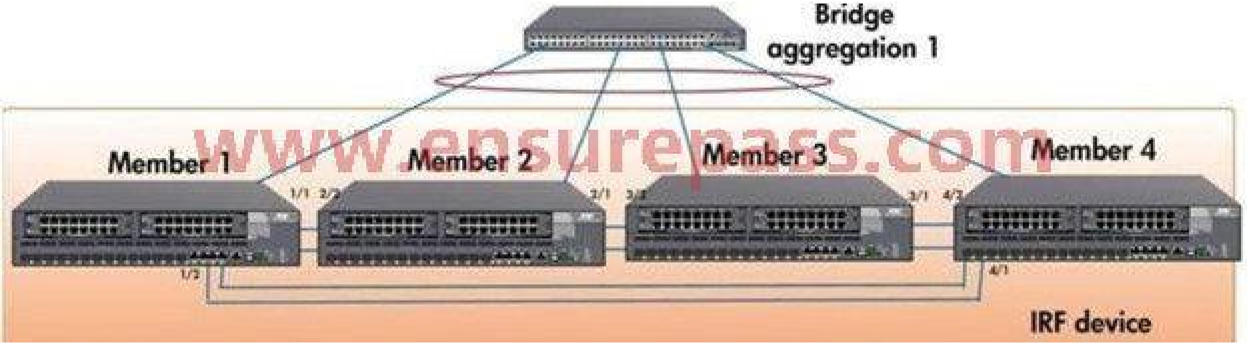

Exhibit 2

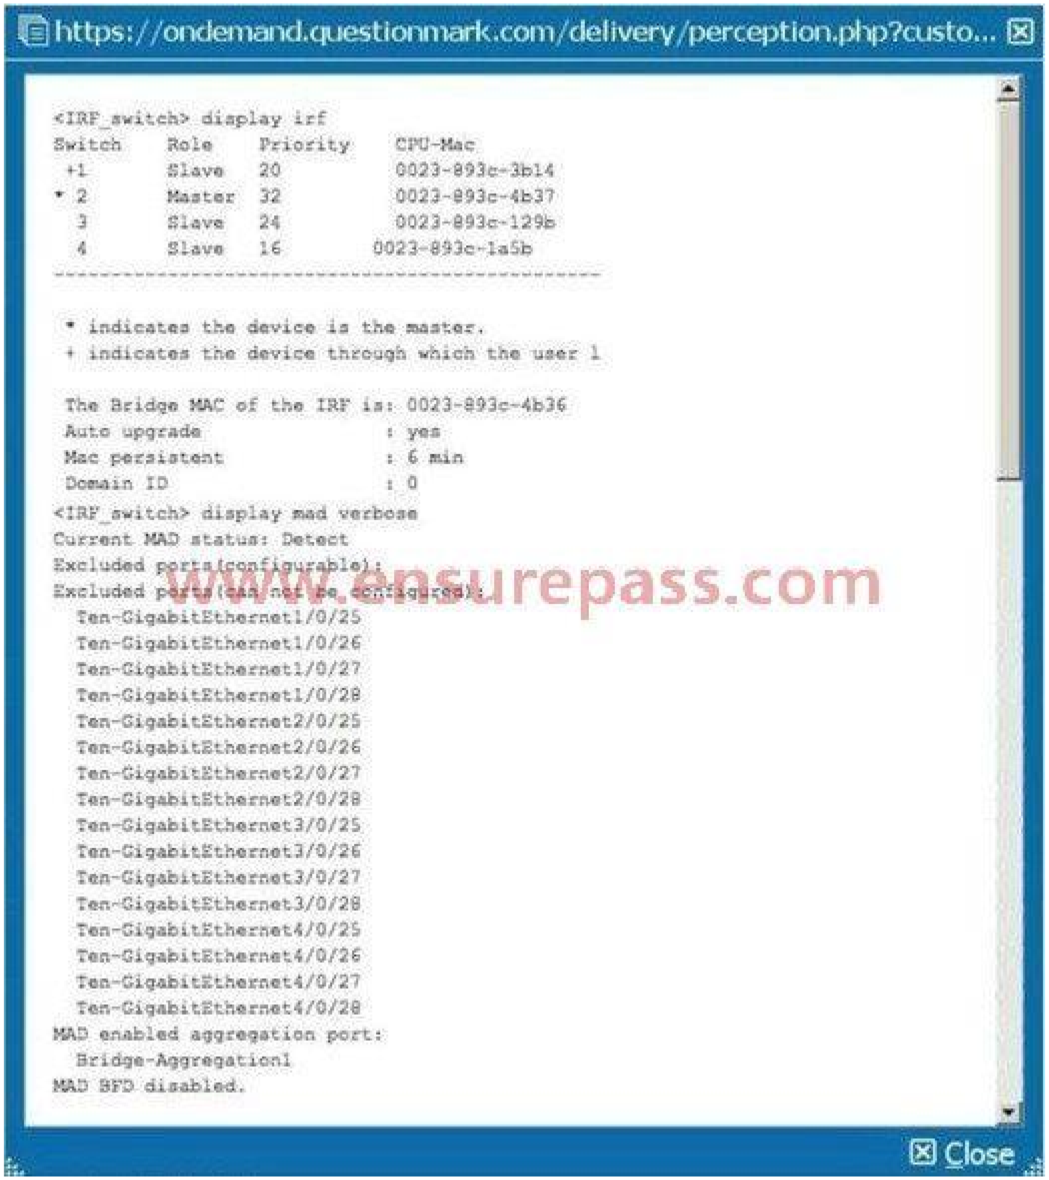
https://ondemand.questionmark.com/delivery/perception.php?custo...
☒ Close

```
<IRF_switch> display irf
Switch     Role    Priority     CPU-Mac
 +1        Slave   20            0023-893c-3b14
* 2        Master  32            0023-893c-4b37
  3        Slave   24            0023-893c-129b
  4        Slave   16          0023-893c-1a5b
-------------------------------------------------

 * indicates the device is the master.
 + indicates the device through which the user l

 The Bridge MAC of the IRF is: 0023-893c-4b36
 Auto upgrade                  : yes
 Mac persistent                : 6 min
 Domain ID                     : 0
<IRF_switch> display mad verbose
Current MAD status: Detect
Excluded ports(configurable):
Excluded ports(can not be configured):
  Ten-GigabitEthernet1/0/25
  Ten-GigabitEthernet1/0/26
  Ten-GigabitEthernet1/0/27
  Ten-GigabitEthernet1/0/28
  Ten-GigabitEthernet2/0/25
  Ten-GigabitEthernet2/0/26
  Ten-GigabitEthernet2/0/27
  Ten-GigabitEthernet2/0/28
  Ten-GigabitEthernet3/0/25
  Ten-GigabitEthernet3/0/26
  Ten-GigabitEthernet3/0/27
  Ten-GigabitEthernet3/0/28
  Ten-GigabitEthernet4/0/25
  Ten-GigabitEthernet4/0/26
  Ten-GigabitEthernet4/0/27
  Ten-GigabitEthernet4/0/28
MAD enabled aggregation port:
  Bridge-Aggregation1
MAD BFD disabled.
```

An IRF device consists of four HP 5820 switches. The IRF topology is shown in Exhibit 1. The IRF settings are shown in Exhibit 2. After these settings are displayed, Member 2 experiences a power failure and shuts down. After an hour, power is restored, and Member 2 reboots. Based on this event and the exhibits, what happens next?

A. Member 2 becomes master, and all other members reboot.
B. Member 2 becomes slave, and all other members reboot.
C. Member 3 remains master, and member 2 reboots.
D. Member 1 remains master, and members 2, 3, and 4 reboot.

Correct Answer: C
Section: High Availability
Explanation

Explanation/Reference:
IRF Role Election
The process of defining the role (master or slave) of members is role election.
Role election is held when the topology changes, such as, forming an IRF virtual device, adding a new member, leaving or failure of the master, or IRF virtual device merge. The master is elected based on the rules below, in the order specified. If the first rule does not apply, a second rule is tried, and so on, until the only winner is found.

- The current master, even if a new member has a higher priority. (When an IRF virtual device is being formed, all member devices consider themselves as the master, so this principle is skipped)
- The device with higher priority.
- The device with the longest system up-time. (The system up-time information of each member device is delivered in IRF hello packets)
- The member with the lowest bridge MAC address.

QUESTION 45
Two HP Provision switches are configured as virtual routers. How does the VRRP backup know if the VRRP master is available and functioning?

A. The VRRP master and backup exchange hello packets over a VLAN dedicated to VRRP.
B. The VRRP backup sends VRRP requests, and the master replies by sending a VRRP ACK message.
C. The VRRP master and backup both implement bi-directional forwarding detection (BFD).
D. The VRRP backup listens for the master's VRRP advertisements.

Correct Answer: D
Section: High Availability
Explanation

Explanation/Reference:

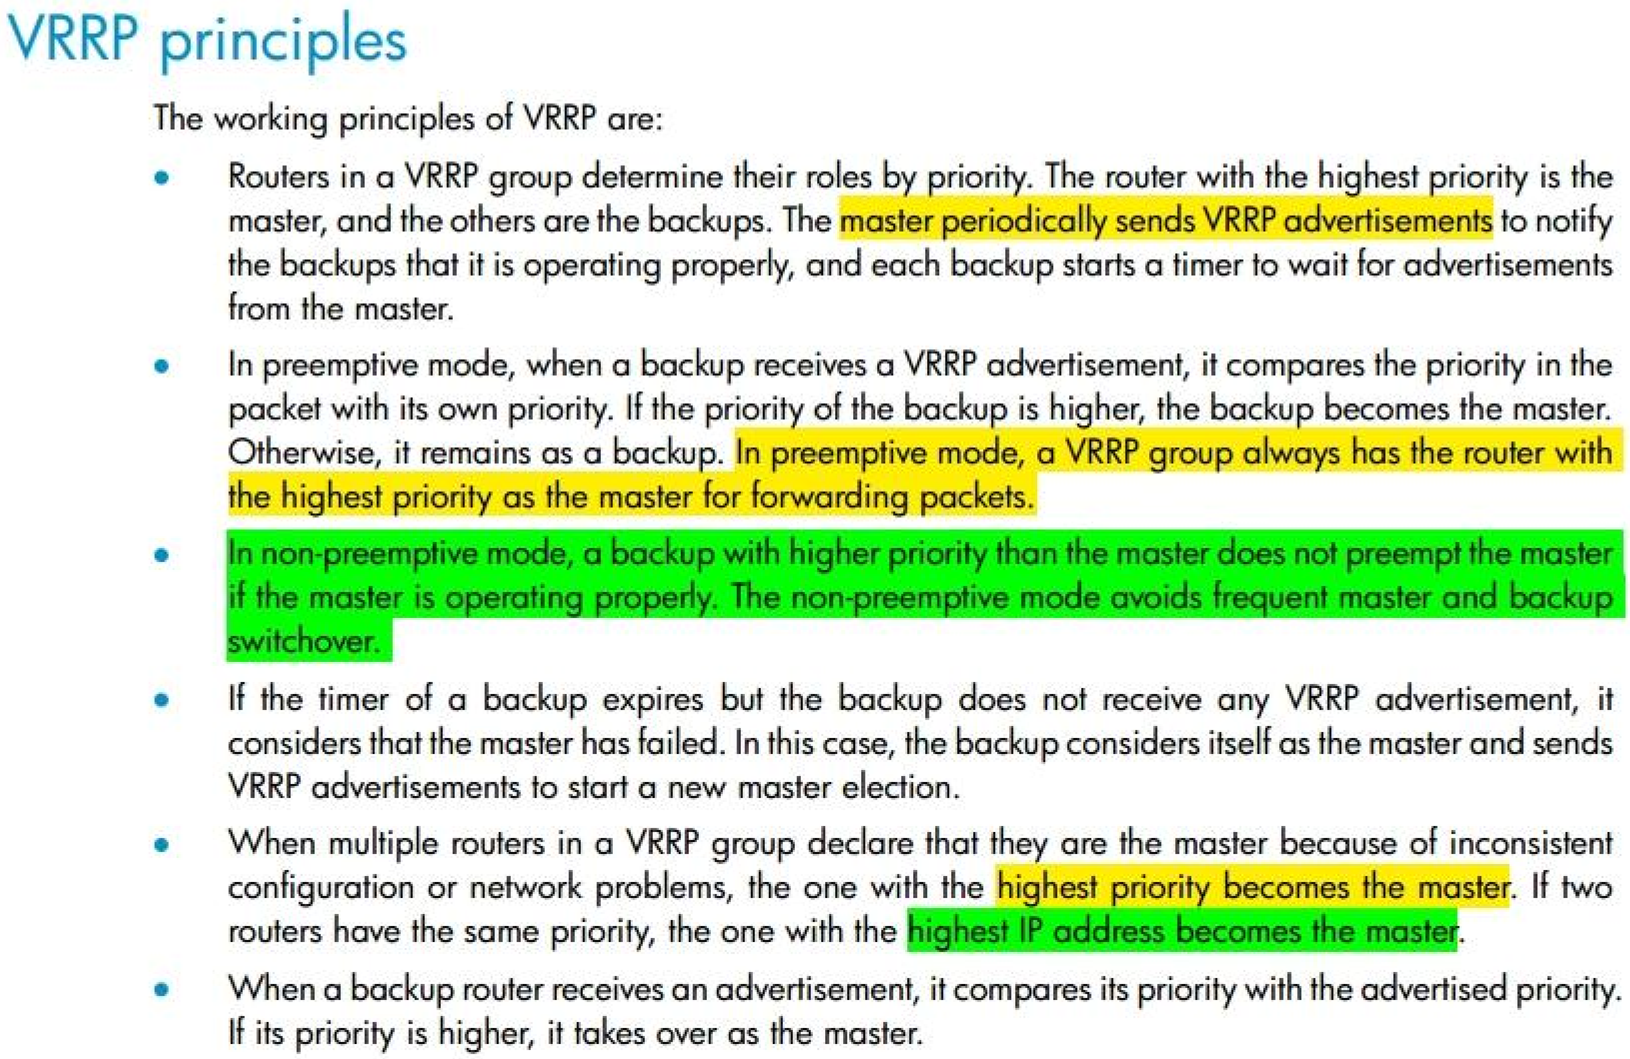

VRRP principles

The working principles of VRRP are:

- Routers in a VRRP group determine their roles by priority. The router with the highest priority is the master, and the others are the backups. The master periodically sends VRRP advertisements to notify the backups that it is operating properly, and each backup starts a timer to wait for advertisements from the master.
- In preemptive mode, when a backup receives a VRRP advertisement, it compares the priority in the packet with its own priority. If the priority of the backup is higher, the backup becomes the master. Otherwise, it remains as a backup. In preemptive mode, a VRRP group always has the router with the highest priority as the master for forwarding packets.
- In non-preemptive mode, a backup with higher priority than the master does not preempt the master if the master is operating properly. The non-preemptive mode avoids frequent master and backup switchover.
- If the timer of a backup expires but the backup does not receive any VRRP advertisement, it considers that the master has failed. In this case, the backup considers itself as the master and sends VRRP advertisements to start a new master election.
- When multiple routers in a VRRP group declare that they are the master because of inconsistent configuration or network problems, the one with the highest priority becomes the master. If two routers have the same priority, the one with the highest IP address becomes the master.
- When a backup router receives an advertisement, it compares its priority with the advertised priority. If its priority is higher, it takes over as the master.

QUESTION 46
View the exhibit. The switches shown in the exhibit are HP Provision ASIC switches. Switches A, B, C, and D implement MSTP using the settings shown in the exhibit. Switches B and C also connect to a switch at the core, switch E, on dedicated VLANs that act like routed links. The network administrator is not certain how or whether switch E implements MSTP. Which action should the administrator take to prevent issues with the connection to the core?

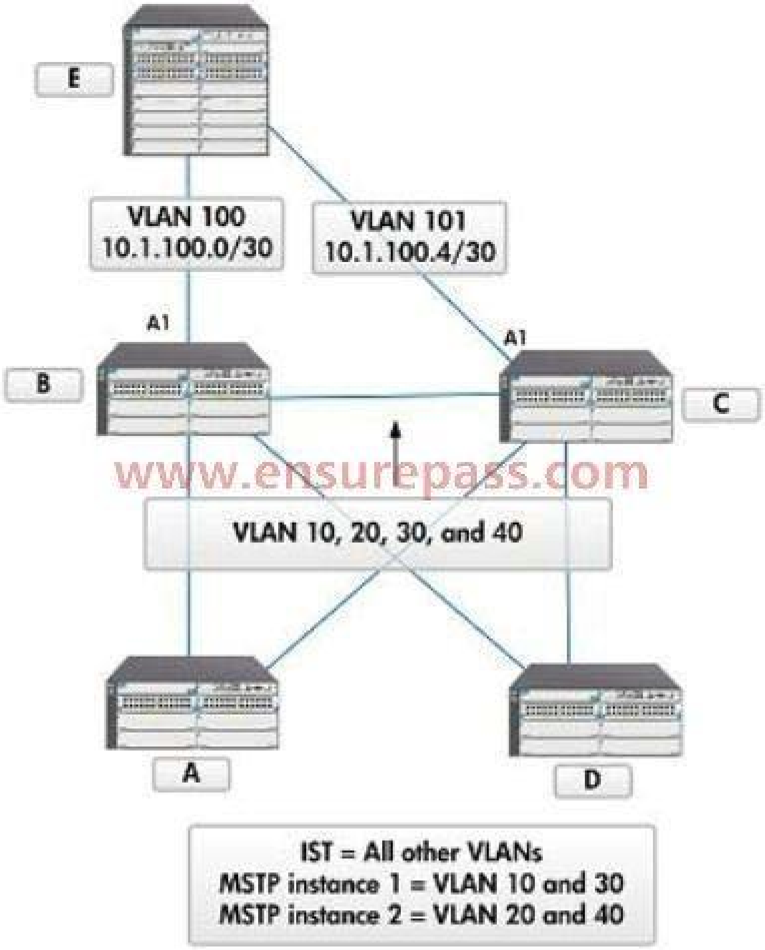

A. Disable MSTP on instance 0.
B. Implement BPDU protection on the A1 ports on switch B and switch C.
C. Implement BPDU filtering on the A1 ports on switch B and switch C.
D. Make VLAN 100 part of MSTP instance 1 and VLAN 101 part of MSTP instance 2.

Correct Answer: C
Section: MSTP
Explanation

Explanation/Reference:
Remember

QUESTION 47
View the exhibit. An HP 5406 zl switch is configured as shown in the exhibit. The switch receives a frame on VLAN 10, destined to UDP port 5555. The frame is marked 802.1p value 6 and DiffServ value 46. Note that queues are numbered 1-8, not 0-7. In which queue does the switch forward the traffic?

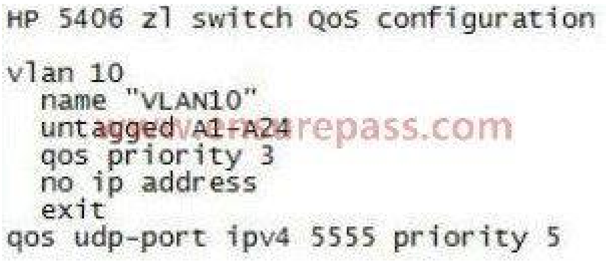

```
HP 5406 zl switch QoS configuration

vlan 10
   name "VLAN10"
   untagged A1-A24
   qos priority 3
   no ip address
   exit
qos udp-port ipv4 5555 priority 5
```

A. 8
B. 4
C. 6
D. 7

Correct Answer: C
Section: QOS
Explanation

Explanation/Reference:
HP ProCurve Switch Software Advanced Traffic Management Guide
...

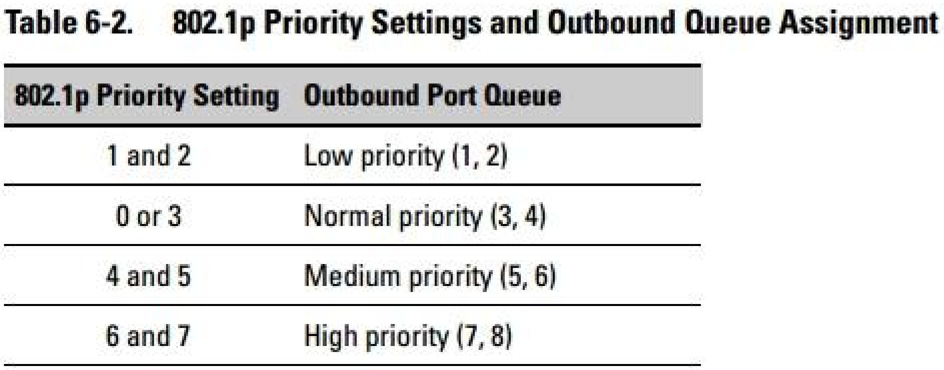

Table 6-2. 802.1p Priority Settings and Outbound Queue Assignment

802.1p Priority Setting	Outbound Port Queue
1 and 2	Low priority (1, 2)
0 or 3	Normal priority (3, 4)
4 and 5	Medium priority (5, 6)
6 and 7	High priority (7, 8)

...

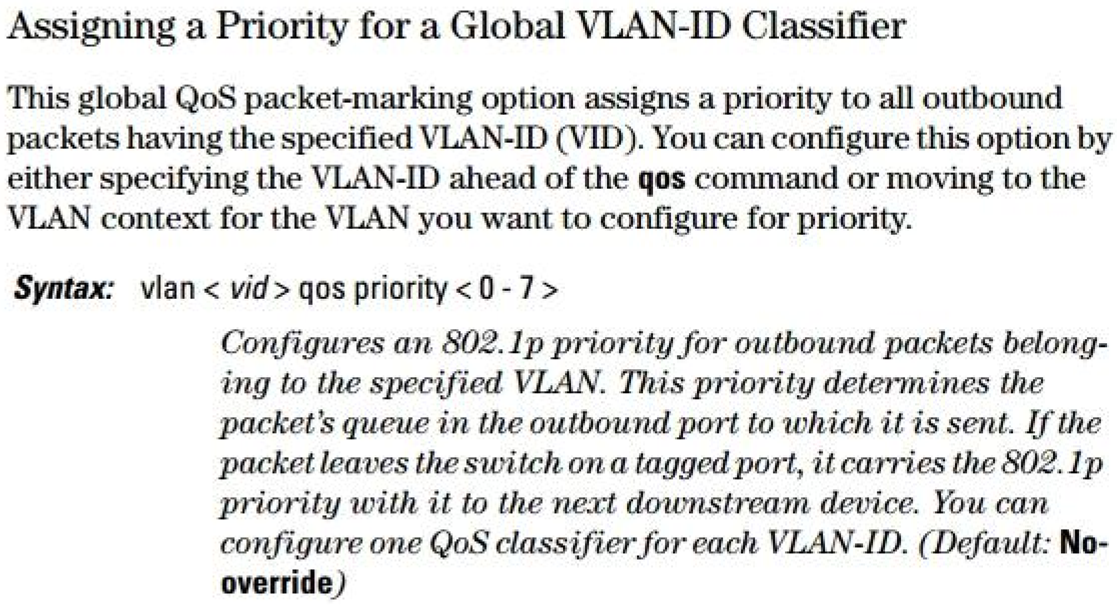

Assigning a Priority for a Global VLAN-ID Classifier

This global QoS packet-marking option assigns a priority to all outbound packets having the specified VLAN-ID (VID). You can configure this option by either specifying the VLAN-ID ahead of the **qos** command or moving to the VLAN context for the VLAN you want to configure for priority.

Syntax: vlan < *vid* > qos priority < 0 - 7 >

Configures an 802.1p priority for outbound packets belonging to the specified VLAN. This priority determines the packet's queue in the outbound port to which it is sent. If the packet leaves the switch on a tagged port, it carries the 802.1p priority with it to the next downstream device. You can configure one QoS classifier for each VLAN-ID. (Default: **No-override***)*

...

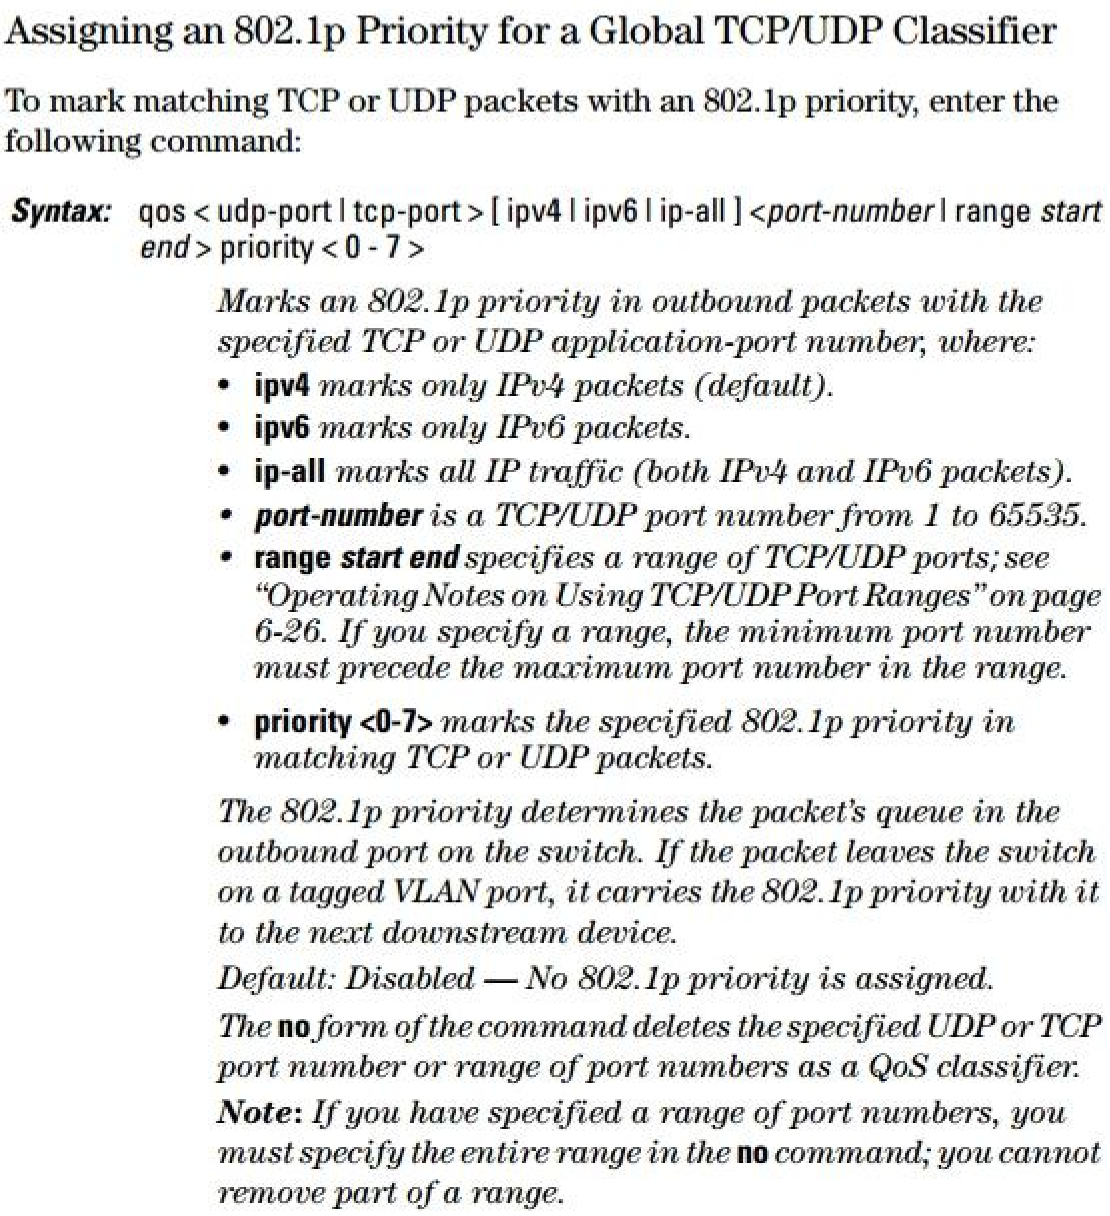

Assigning an 802.1p Priority for a Global TCP/UDP Classifier

To mark matching TCP or UDP packets with an 802.1p priority, enter the following command:

Syntax: qos < udp-port | tcp-port > [ipv4 | ipv6 | ip-all] <*port-number* | range *start end* > priority < 0 - 7 >

Marks an 802.1p priority in outbound packets with the specified TCP or UDP application-port number, where:

- **ipv4** *marks only IPv4 packets (default).*
- **ipv6** *marks only IPv6 packets.*
- **ip-all** *marks all IP traffic (both IPv4 and IPv6 packets).*
- ***port-number*** *is a TCP/UDP port number from 1 to 65535.*
- **range *start end*** *specifies a range of TCP/UDP ports; see "Operating Notes on Using TCP/UDP Port Ranges" on page 6-26. If you specify a range, the minimum port number must precede the maximum port number in the range.*
- **priority <0-7>** *marks the specified 802.1p priority in matching TCP or UDP packets.*

The 802.1p priority determines the packet's queue in the outbound port on the switch. If the packet leaves the switch on a tagged VLAN port, it carries the 802.1p priority with it to the next downstream device.

Default: Disabled — No 802.1p priority is assigned.

The **no** *form of the command deletes the specified UDP or TCP port number or range of port numbers as a QoS classifier.*

Note: *If you have specified a range of port numbers, you must specify the entire range in the* **no** *command; you cannot remove part of a range.*

...

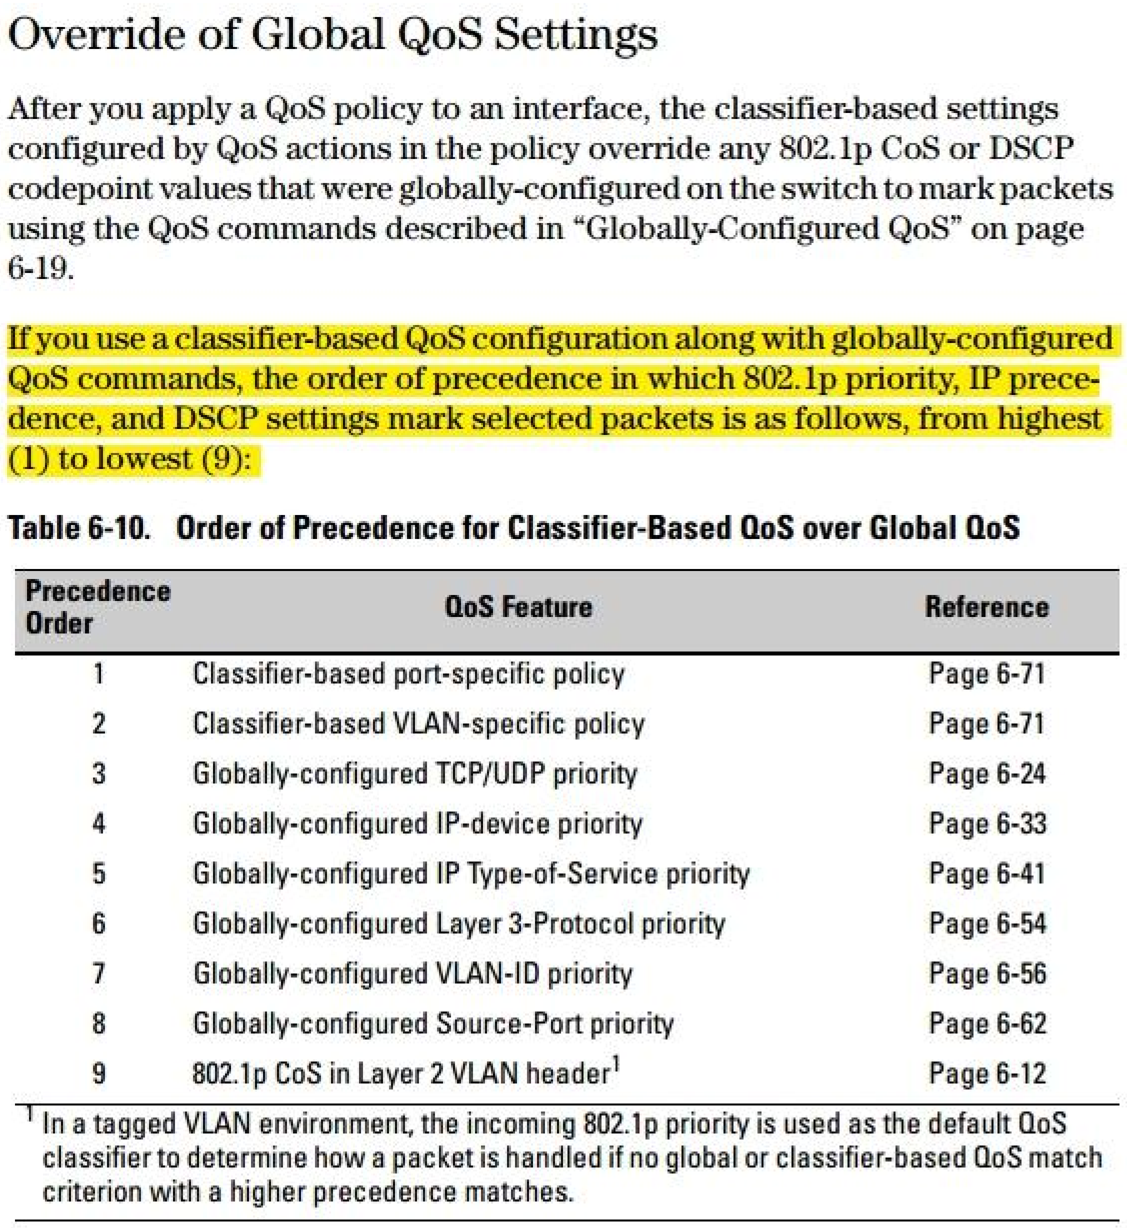

Override of Global QoS Settings

After you apply a QoS policy to an interface, the classifier-based settings configured by QoS actions in the policy override any 802.1p CoS or DSCP codepoint values that were globally-configured on the switch to mark packets using the QoS commands described in “Globally-Configured QoS” on page 6-19.

If you use a classifier-based QoS configuration along with globally-configured QoS commands, the order of precedence in which 802.1p priority, IP precedence, and DSCP settings mark selected packets is as follows, from highest (1) to lowest (9):

Table 6-10. Order of Precedence for Classifier-Based QoS over Global QoS

Precedence Order	QoS Feature	Reference
1	Classifier-based port-specific policy	Page 6-71
2	Classifier-based VLAN-specific policy	Page 6-71
3	Globally-configured TCP/UDP priority	Page 6-24
4	Globally-configured IP-device priority	Page 6-33
5	Globally-configured IP Type-of-Service priority	Page 6-41
6	Globally-configured Layer 3-Protocol priority	Page 6-54
7	Globally-configured VLAN-ID priority	Page 6-56
8	Globally-configured Source-Port priority	Page 6-62
9	802.1p CoS in Layer 2 VLAN header[1]	Page 6-12

[1] In a tagged VLAN environment, the incoming 802.1p priority is used as the default QoS classifier to determine how a packet is handled if no global or classifier-based QoS match criterion with a higher precedence matches.

QUESTION 48
A company wants clients to receive a video stream from a server. What is the advantage of transmitting this stream as multicast traffic rather than unicast traffic?

A. The network administrator does not need to configure routing for the stream.
B. The server does not need to be preconfigured with each endpoint's IP address.
C. The endpoints can reside in many different subnets instead of only the server's local subnet.
D. The server transmits a single stream rather than many streams.

Correct Answer: D
Section: Multicast
Explanation

Explanation/Reference:
Remember

QUESTION 49
View the exhibit. A network administrator configures the HP 2610-24-PoE switch as shown in the exhibit. Computers belong to the data VLAN, which is VLAN 10. IP phones should transmit traffic on VLAN 20. They should be able to connect to any port and use LLDP-MED to obtain their VLAN assignment. However, IP phones connected to this switch cannot reach the VoIP server and are not functioning correctly. IP phones connected to other switches are functioning correctly. What must the network administrator do?

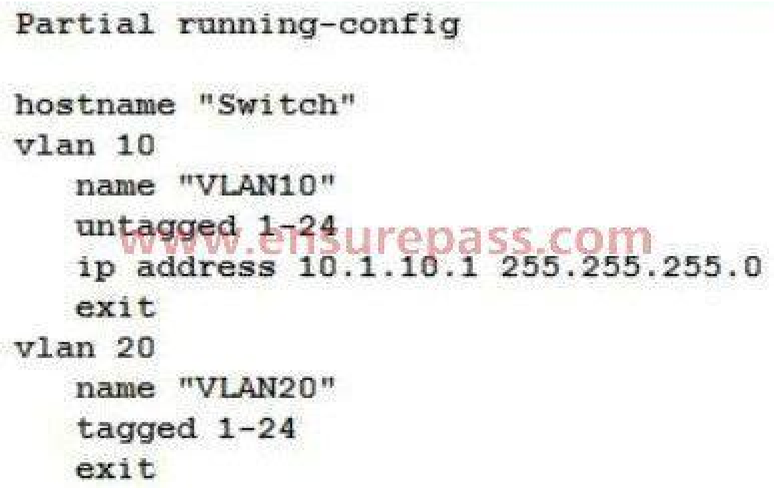

```
Partial running-config

hostname "Switch"
vlan 10
   name "VLAN10"
   untagged 1-24
   ip address 10.1.10.1 255.255.255.0
   exit
vlan 20
   name "VLAN20"
   tagged 1-24
   exit
```

A. Make VLAN 20 untagged on the ports into which phones are connected.
B. Define VLAN 20 as a voice VLAN.

C. Enable LLDP-MED globally.

D. Enable LLDP-MED on the ports into which phones are connected.

Correct Answer: B
Section: Other
Explanation

Explanation/Reference:
Remember
LLDP-MED enabled globally by default!

QUESTION 50
View the exhibit. Several HP Provision ASIC switches are implementing PIM. Based on the configurations shown in the exhibit, which switch is included in the dynamic RP set?

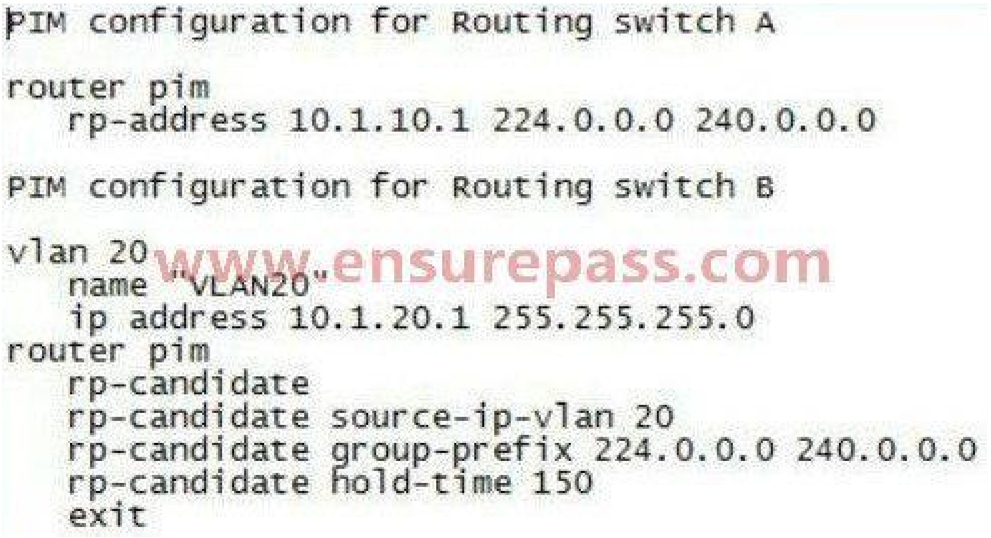

```
PIM configuration for Routing switch A

router pim
   rp-address 10.1.10.1 224.0.0.0 240.0.0.0

PIM configuration for Routing switch B

vlan 20
   name "VLAN20"
   ip address 10.1.20.1 255.255.255.0
router pim
   rp-candidate
   rp-candidate source-ip-vlan 20
   rp-candidate group-prefix 224.0.0.0 240.0.0.0
   rp-candidate hold-time 150
   exit
```

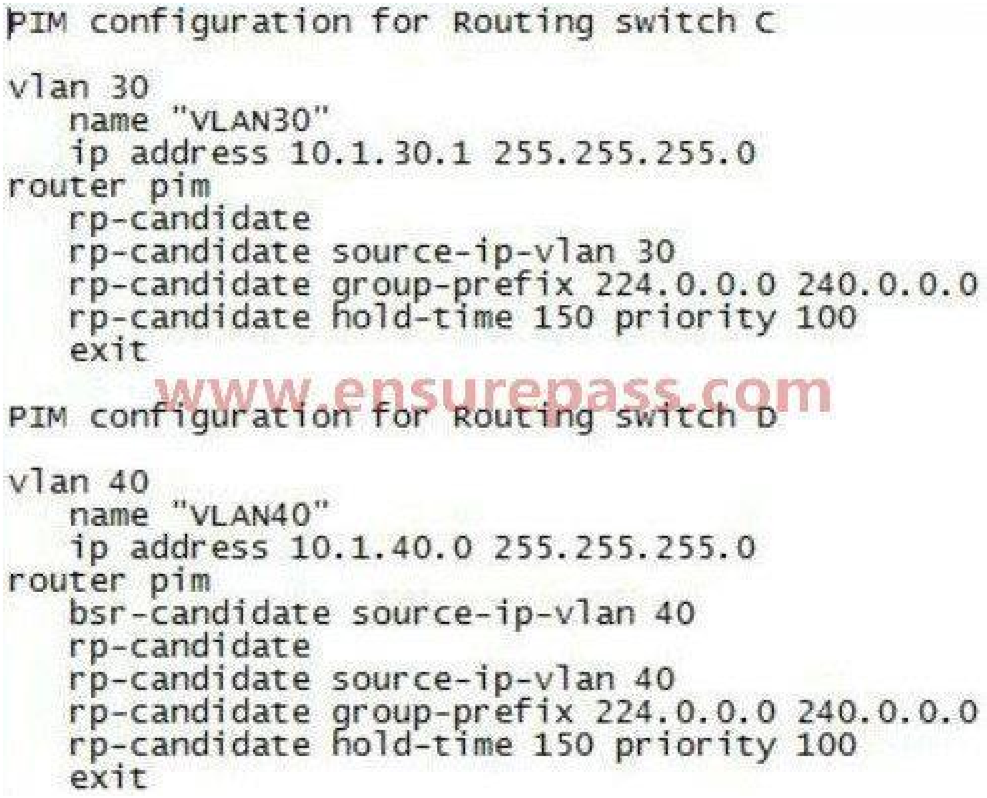

```
PIM configuration for Routing switch C

vlan 30
   name "VLAN30"
   ip address 10.1.30.1 255.255.255.0
router pim
   rp-candidate
   rp-candidate source-ip-vlan 30
   rp-candidate group-prefix 224.0.0.0 240.0.0.0
   rp-candidate hold-time 150 priority 100
   exit

PIM configuration for Routing switch D

vlan 40
   name "VLAN40"
   ip address 10.1.40.0 255.255.255.0
router pim
   bsr-candidate source-ip-vlan 40
   rp-candidate
   rp-candidate source-ip-vlan 40
   rp-candidate group-prefix 224.0.0.0 240.0.0.0
   rp-candidate hold-time 150 priority 100
   exit
```

A. routing switch A
B. routing switch B
C. routing switch D
D. routing switch C

Correct Answer: D
Section: Multicast
Explanation

Explanation/Reference:
On the switch 'A' is configured not the correct IP address, so it does not participate in the elections!

Candidate-RP Election. Within a PIM-SM domain, different RPs support different multicast addresses or ranges of multicast addresses. (That is, a given PIM-SM multicast group or range of groups is supported by only one active RP, although other candidate RPs can also be configured with overlapping or identical support.)

A candidate RP's group-prefix configuration identifies the multicast groups the RP is enabled to support.

If multiple candidate RPs have group prefixes configured so that any of these RPs can support a given multicast group, then the following criteria are used to select the RP to support the group:

1. The C-RP configured with the longest group-prefix mask applicable to the multicast group is selected to support the group. If multiple RP candidates meet this criterion, then step 2 applies.
2. The C-RP configured with the highest priority is selected. If multiple RP candidates meet this criterion, then step 3 applies.
3. A hash function (using the configured **bsr-candidate hash-mask-length** value) generates a series of mask length values that are individually assigned to the set of eligible C-RPs. If the hash function matches a single RP candidate to a longer mask length than the other candidates, that candidate is selected to support the group. If the hash function matches the longest mask length to multiple RP candidates, then step 4 applies.
4. The C-RP having the highest IP address is selected to support the group.

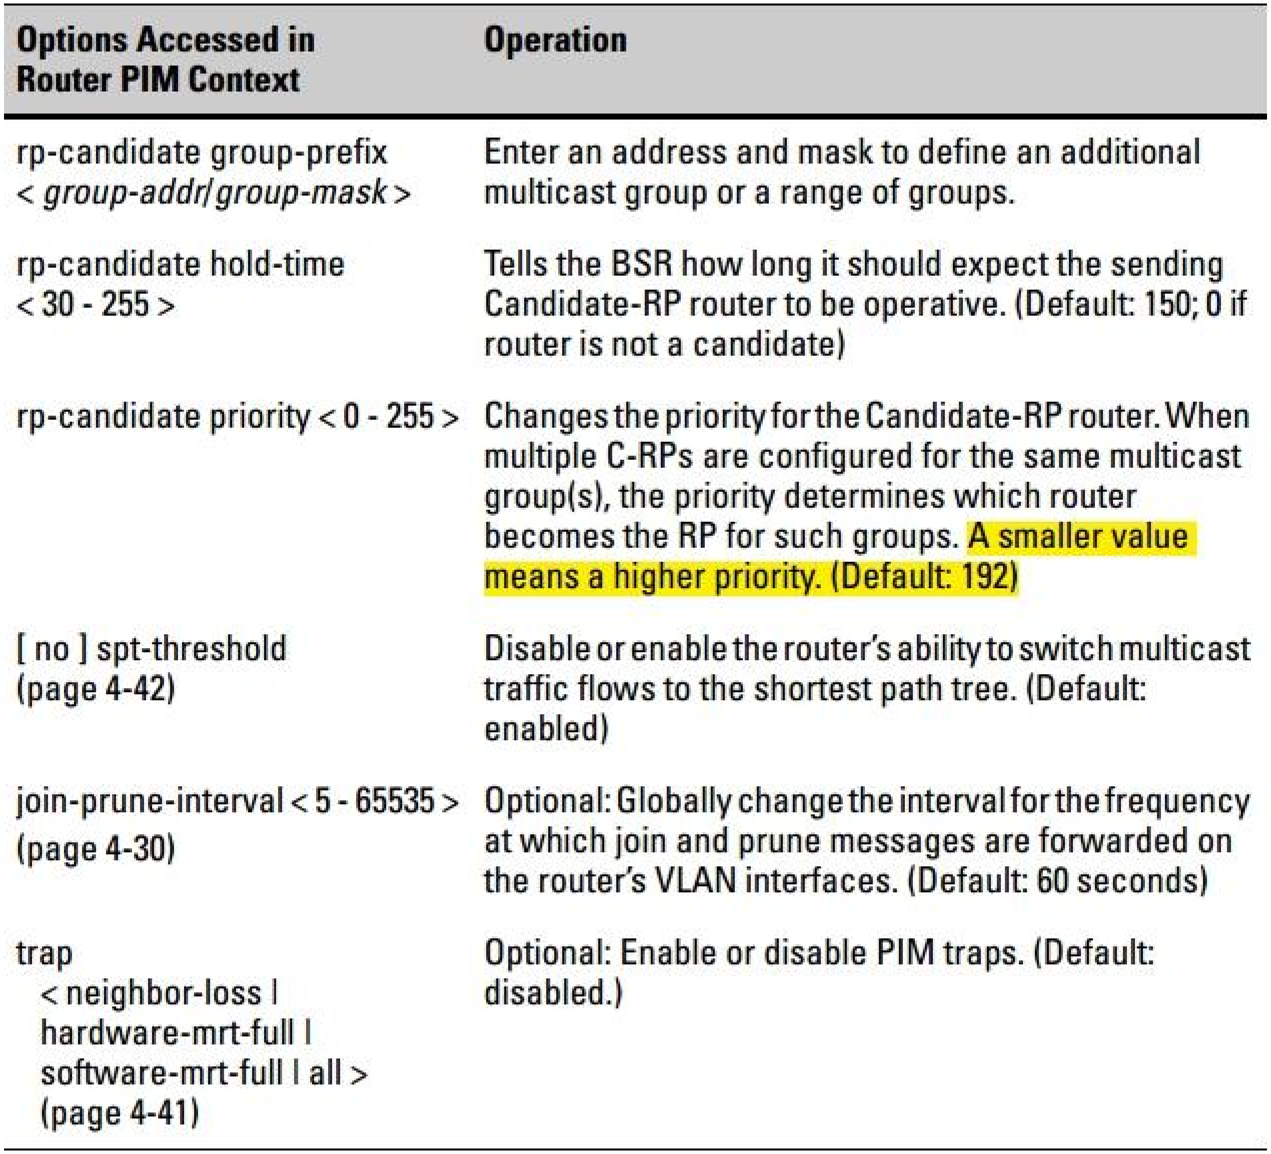

Options Accessed in Router PIM Context	Operation
rp-candidate group-prefix < *group-addr/group-mask* >	Enter an address and mask to define an additional multicast group or a range of groups.
rp-candidate hold-time < 30 - 255 >	Tells the BSR how long it should expect the sending Candidate-RP router to be operative. (Default: 150; 0 if router is not a candidate)
rp-candidate priority < 0 - 255 >	Changes the priority for the Candidate-RP router. When multiple C-RPs are configured for the same multicast group(s), the priority determines which router becomes the RP for such groups. A smaller value means a higher priority. (Default: 192)
[no] spt-threshold (page 4-42)	Disable or enable the router's ability to switch multicast traffic flows to the shortest path tree. (Default: enabled)
join-prune-interval < 5 - 65535 > (page 4-30)	Optional: Globally change the interval for the frequency at which join and prune messages are forwarded on the router's VLAN interfaces. (Default: 60 seconds)
trap < neighbor-loss | hardware-mrt-full | software-mrt-full | all > (page 4-41)	Optional: Enable or disable PIM traps. (Default: disabled.)

QUESTION 51

View the exhibit.

Exhibit 1

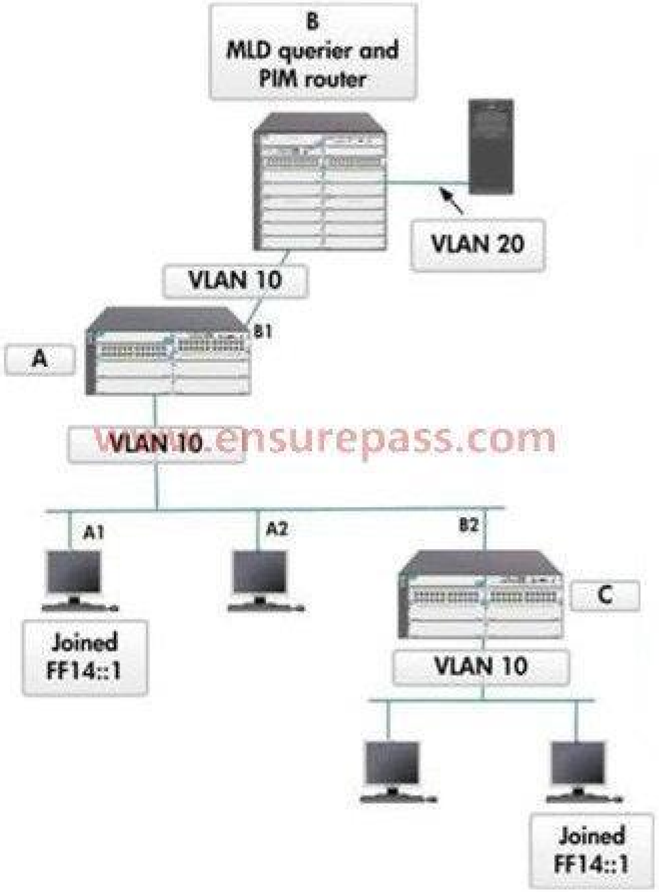

Exhibit 2

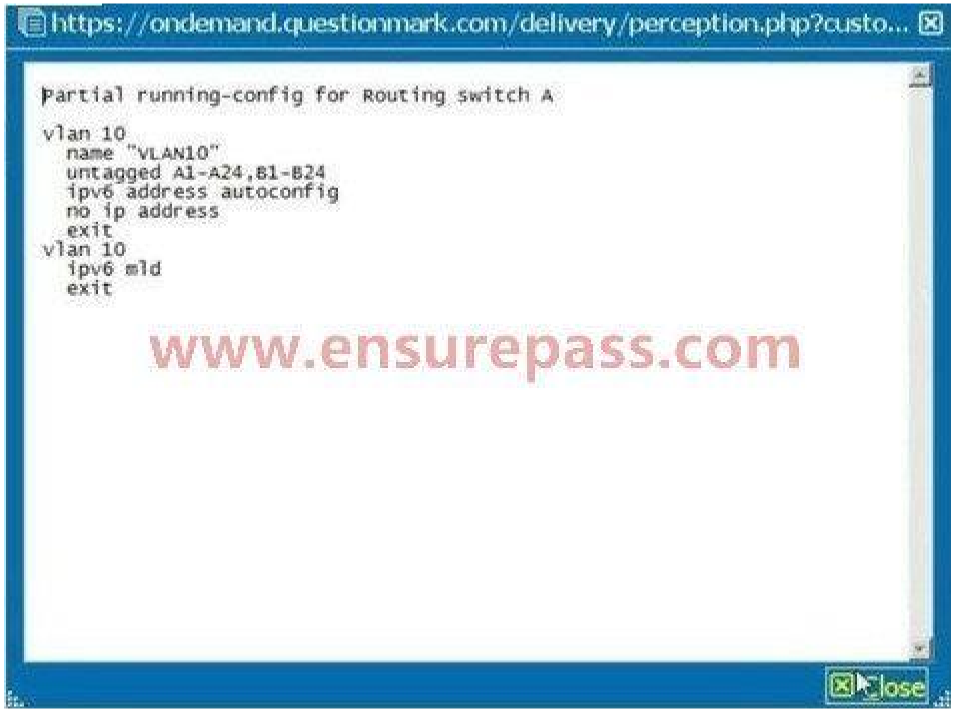

https://ondemand.questionmark.com/delivery/perception.php?custo...

```
Partial running-config for Routing switch A

vlan 10
   name "VLAN10"
   untagged A1-A24,B1-B24
   ipv6 address autoconfig
   no ip address
   exit
vlan 10
   ipv6 mld
   exit
```

A network that handles IPv6 multicasts and the hosts that have joined particular IPv6 multicast groups is shown in Exhibit 1. Routing switch B is acting as the MLD querier, and routing switch A implements MLD snooping, as shown in Exhibit 2. An IPv6 multicast packet for FF14::1 arrives on routing switch A on port B1. On which ports does the switch forward the packet?

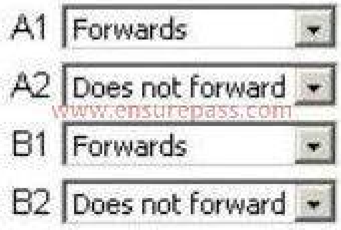

A1 Forwards
A2 Does not forward
B1 Forwards
B2 Does not forward

A. A1 Forwards - A2 Does not forward - B1 Does not forward - B2 Forwards
B. A1 Does not forward - A2 Forwards - B1 Forwards - B2 Does not forward
C. A1 Forwards - A2 Does not forward - B1 Forwards - B2 Does not forward
D. A1 Does not forward - A2 Forwards - B1 Does not forward - B2 Forwards

Correct Answer: A
Section: Multicast
Explanation

Explanation/Reference:

General operation. Multicast communication can take place without MLD, and by default MLD is disabled. In that case, if a switch receives a packet with a multicast destination address, it floods the packet to all ports in the same VLAN (except the port that it came in on). Any network nodes that are listening to that multicast address will see the packet; all other hosts ignore the packet.

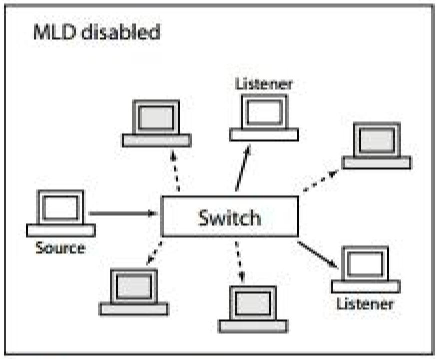

Figure 7-1. Without MLD, multicast traffic is flooded to all ports.

When MLD snooping is enabled on a VLAN, the switch acts to minimize unnecessary multicast traffic. If the switch receives multicast traffic destined for a given multicast address, it forwards that traffic only to ports on the VLAN that have MLD hosts for that address. It drops that traffic for ports on the VLAN that have no MLD hosts (except for a few special cases explained below).

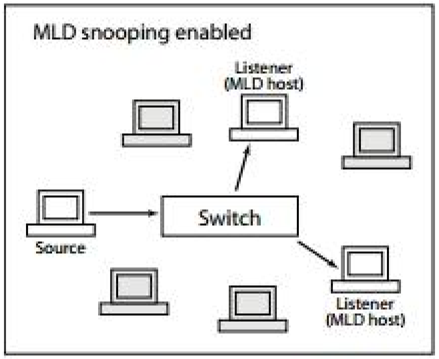

Figure 7-2. With MLD snooping, traffic is sent to MLD hosts.

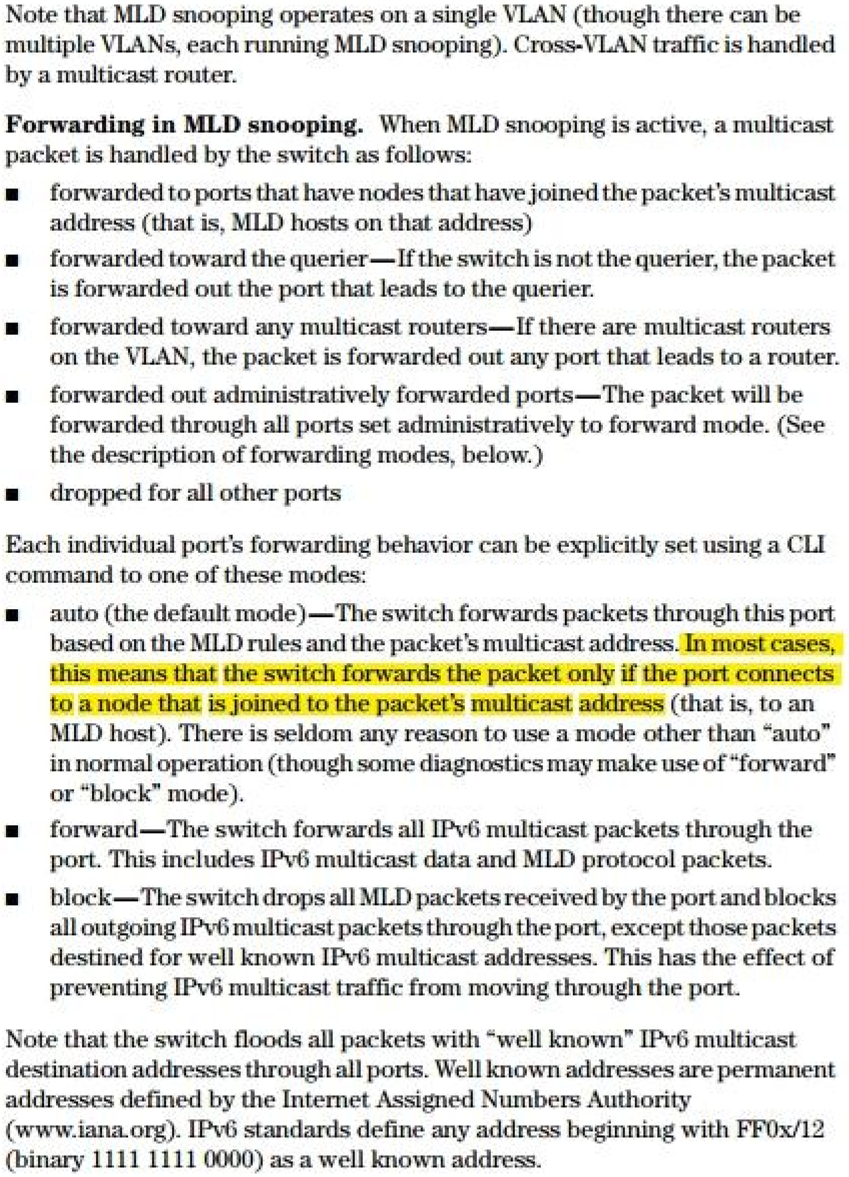

Note that MLD snooping operates on a single VLAN (though there can be multiple VLANs, each running MLD snooping). Cross-VLAN traffic is handled by a multicast router.

Forwarding in MLD snooping. When MLD snooping is active, a multicast packet is handled by the switch as follows:

- forwarded to ports that have nodes that have joined the packet's multicast address (that is, MLD hosts on that address)
- forwarded toward the querier—If the switch is not the querier, the packet is forwarded out the port that leads to the querier.
- forwarded toward any multicast routers—If there are multicast routers on the VLAN, the packet is forwarded out any port that leads to a router.
- forwarded out administratively forwarded ports—The packet will be forwarded through all ports set administratively to forward mode. (See the description of forwarding modes, below.)
- dropped for all other ports

Each individual port's forwarding behavior can be explicitly set using a CLI command to one of these modes:

- auto (the default mode)—The switch forwards packets through this port based on the MLD rules and the packet's multicast address. In most cases, this means that the switch forwards the packet only if the port connects to a node that is joined to the packet's multicast address (that is, to an MLD host). There is seldom any reason to use a mode other than "auto" in normal operation (though some diagnostics may make use of "forward" or "block" mode).
- forward—The switch forwards all IPv6 multicast packets through the port. This includes IPv6 multicast data and MLD protocol packets.
- block—The switch drops all MLD packets received by the port and blocks all outgoing IPv6 multicast packets through the port, except those packets destined for well known IPv6 multicast addresses. This has the effect of preventing IPv6 multicast traffic from moving through the port.

Note that the switch floods all packets with "well known" IPv6 multicast destination addresses through all ports. Well known addresses are permanent addresses defined by the Internet Assigned Numbers Authority (www.iana.org). IPv6 standards define any address beginning with FF0x/12 (binary 1111 1111 0000) as a well known address.

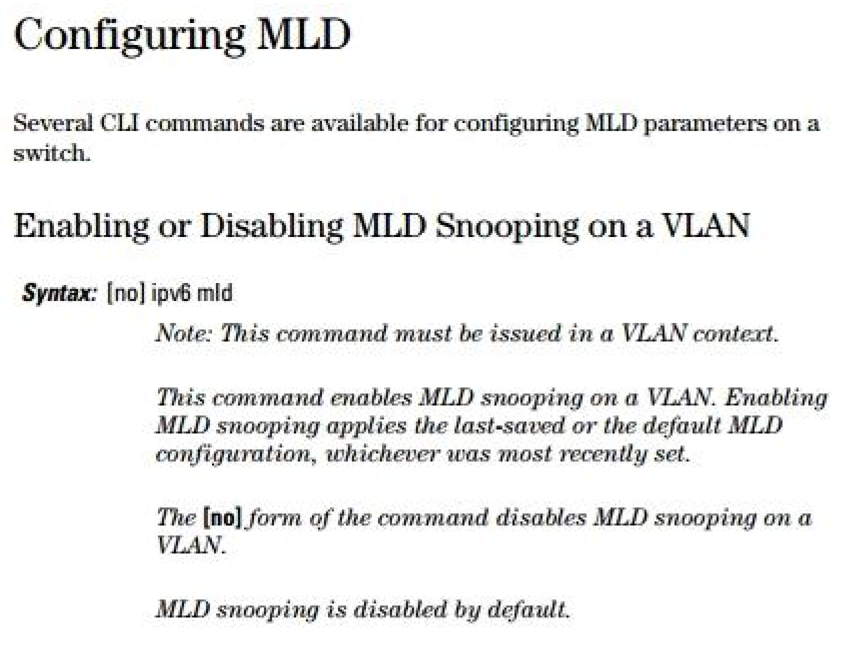

Configuring MLD

Several CLI commands are available for configuring MLD parameters on a switch.

Enabling or Disabling MLD Snooping on a VLAN

Syntax: [no] ipv6 mld

Note: This command must be issued in a VLAN context.

This command enables MLD snooping on a VLAN. Enabling MLD snooping applies the last-saved or the default MLD configuration, whichever was most recently set.

The **[no]** *form of the command disables MLD snooping on a VLAN.*

MLD snooping is disabled by default.

QUESTION 52
A company's HP Provision ASIC switches are configured to use PIM-SM mode to provide multicast services. When an HP Provision ASIC switch receives a multicast stream, it determines that the best path to the multicast source is not through the Rendezvous Point (RP). Which action might this switch take?

A. The switch may use the backup RP if this RP is in the path between the multicast source and its receivers.
B. The switch may send a graft message on the interface through which it reaches the multicast source.
C. The switch may establish itself as the new RP for the shortest path tree.
D. The switch may decide to join the shortest path tree.

Correct Answer: D
Section: Multicast
Explanation

Explanation/Reference:

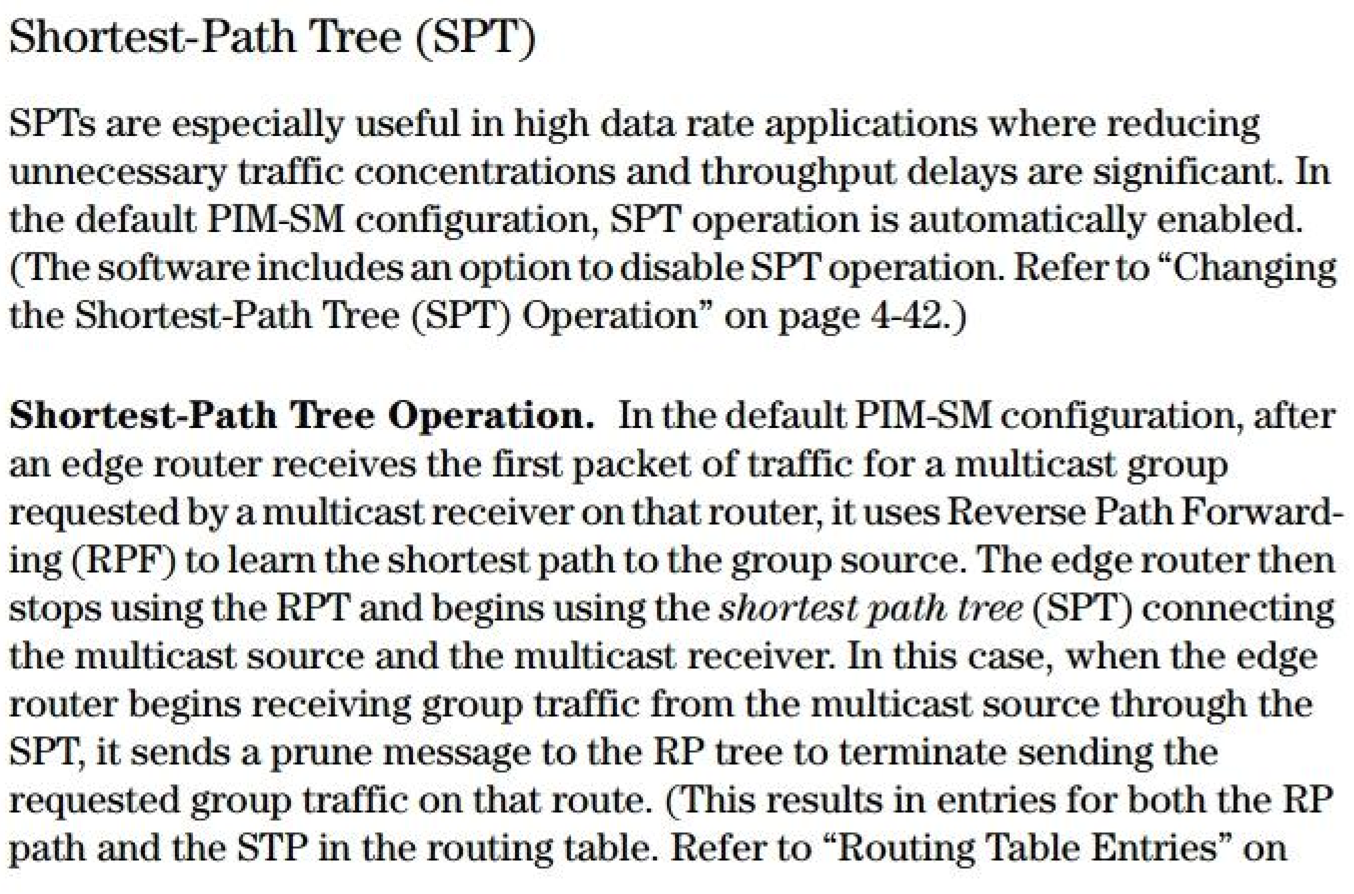

Shortest-Path Tree (SPT)

SPTs are especially useful in high data rate applications where reducing unnecessary traffic concentrations and throughput delays are significant. In the default PIM-SM configuration, SPT operation is automatically enabled. (The software includes an option to disable SPT operation. Refer to “Changing the Shortest-Path Tree (SPT) Operation” on page 4-42.)

Shortest-Path Tree Operation. In the default PIM-SM configuration, after an edge router receives the first packet of traffic for a multicast group requested by a multicast receiver on that router, it uses Reverse Path Forwarding (RPF) to learn the shortest path to the group source. The edge router then stops using the RPT and begins using the *shortest path tree* (SPT) connecting the multicast source and the multicast receiver. In this case, when the edge router begins receiving group traffic from the multicast source through the SPT, it sends a prune message to the RP tree to terminate sending the requested group traffic on that route. (This results in entries for both the RP path and the STP in the routing table. Refer to “Routing Table Entries” on page 4-67.) When completed, the switchover from the RPT to a shorter SPT can reduce unnecessary traffic concentrations in the network and reduce multicast traffic throughput delays.

Note that the switchover from RPT to SPT is not instantaneous. For a short period, packets for a given multicast group may be received from both the RPT and the SPT. Also, in some topologies, the RPT and the SPT to the same edge router may be identical.

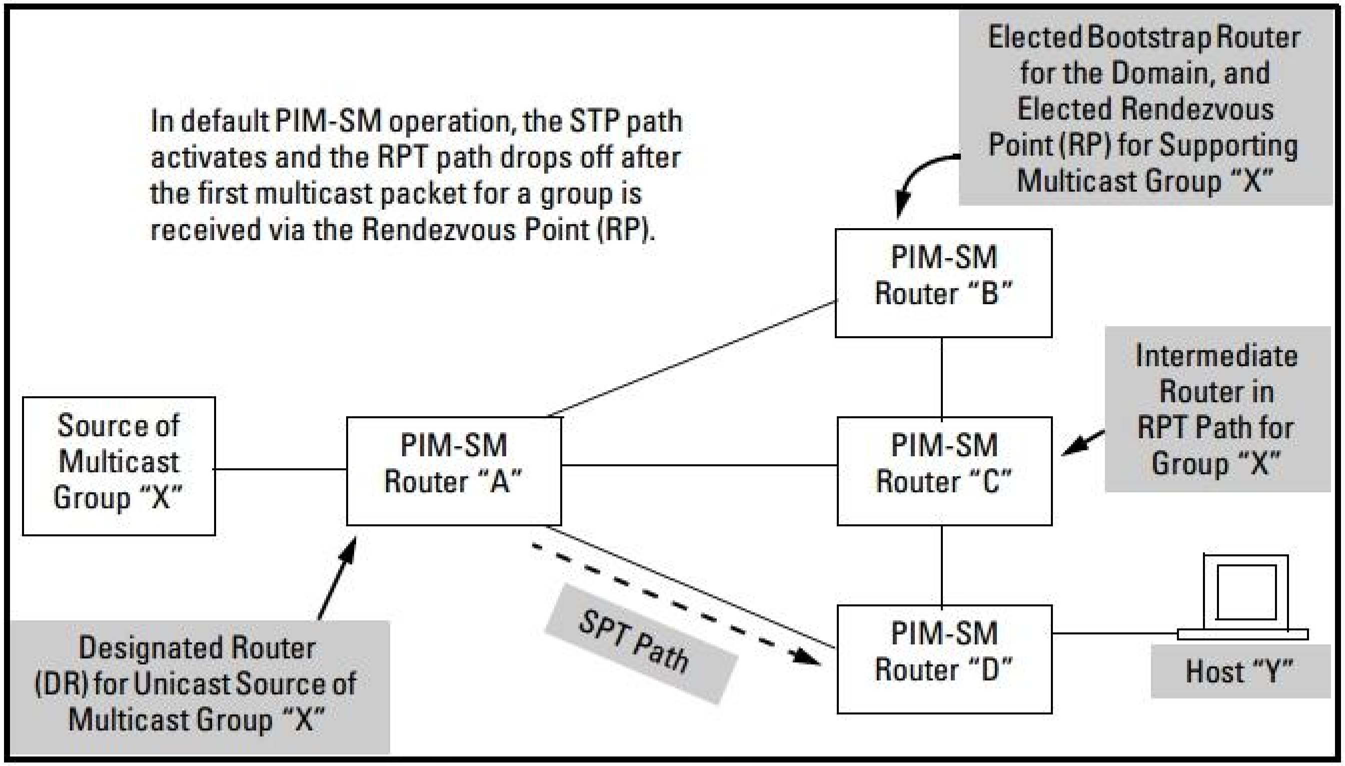
In default PIM-SM operation, the STP path activates and the RPT path drops off after the first multicast packet for a group is received via the Rendezvous Point (RP).
Elected Bootstrap Router for the Domain, and Elected Rendezvous Point (RP) for Supporting Multicast Group "X"
PIM-SM Router "B"
Intermediate Router in RPT Path for Group "X"
Source of Multicast Group "X"
PIM-SM Router "A"
PIM-SM Router "C"
PIM-SM Router "D"
Host "Y"
SPT Path
Designated Router (DR) for Unicast Source of Multicast Group "X"

Figure 4-2. Example PIM-SM Domain with SPT Active To Support a Host that Has Joined a Multicast Group

QUESTION 53
If you use the MAC lockout feature to block a specific MAC address on an HP 3500zl switch, which traffic is dropped?

A. The switch will drop traffic from devices directly connected to the specific port on which MAC lockout is enabled if the destination or source address is the specified MAC address.
B. The switch will drop frames only if the source address is the specified MAC address and port security is also configured on the receiving port.
C. The switch will drop any Layer 2 traffic that contains the specified MAC address as the source address.
D. The switch will drop routed or switched traffic if the destination is the specified MAC address.

Correct Answer: C
Section: Security
Explanation

Explanation/Reference:

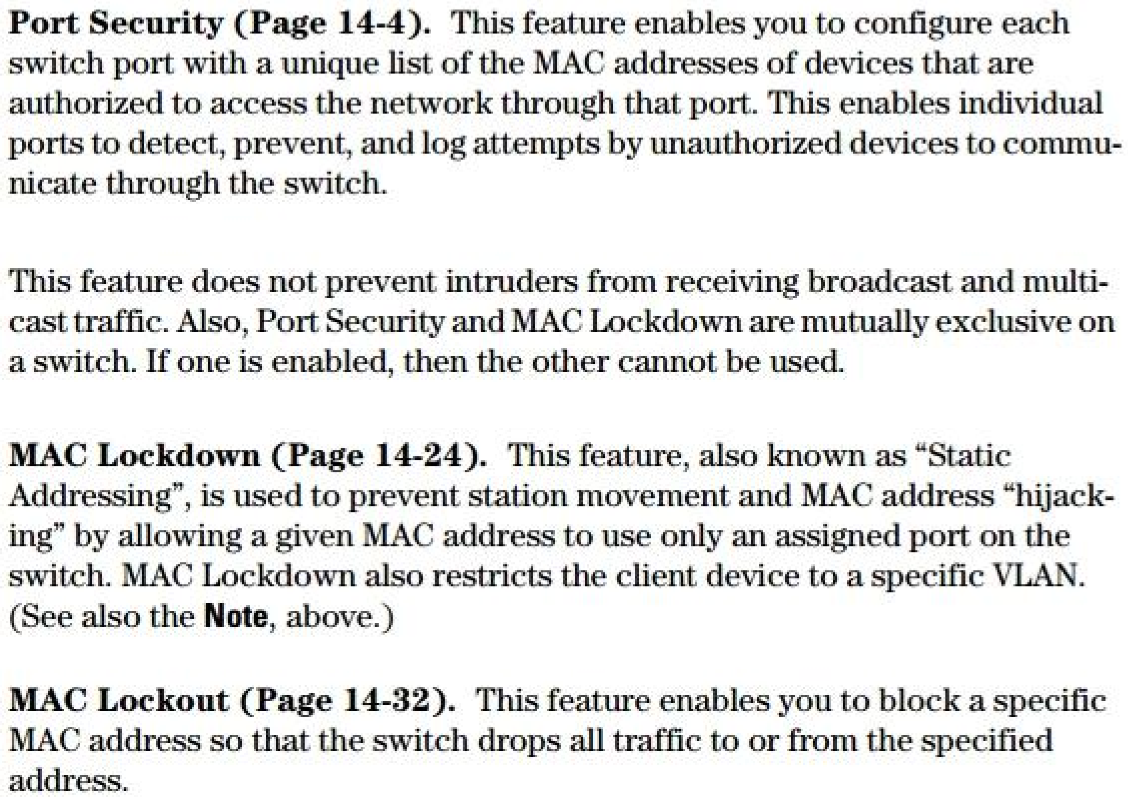

Port Security (Page 14-4). This feature enables you to configure each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch.

This feature does not prevent intruders from receiving broadcast and multicast traffic. Also, Port Security and MAC Lockdown are mutually exclusive on a switch. If one is enabled, then the other cannot be used.

MAC Lockdown (Page 14-24). This feature, also known as “Static Addressing”, is used to prevent station movement and MAC address “hijacking” by allowing a given MAC address to use only an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. (See also the **Note**, above.)

MAC Lockout (Page 14-32). This feature enables you to block a specific MAC address so that the switch drops all traffic to or from the specified address.

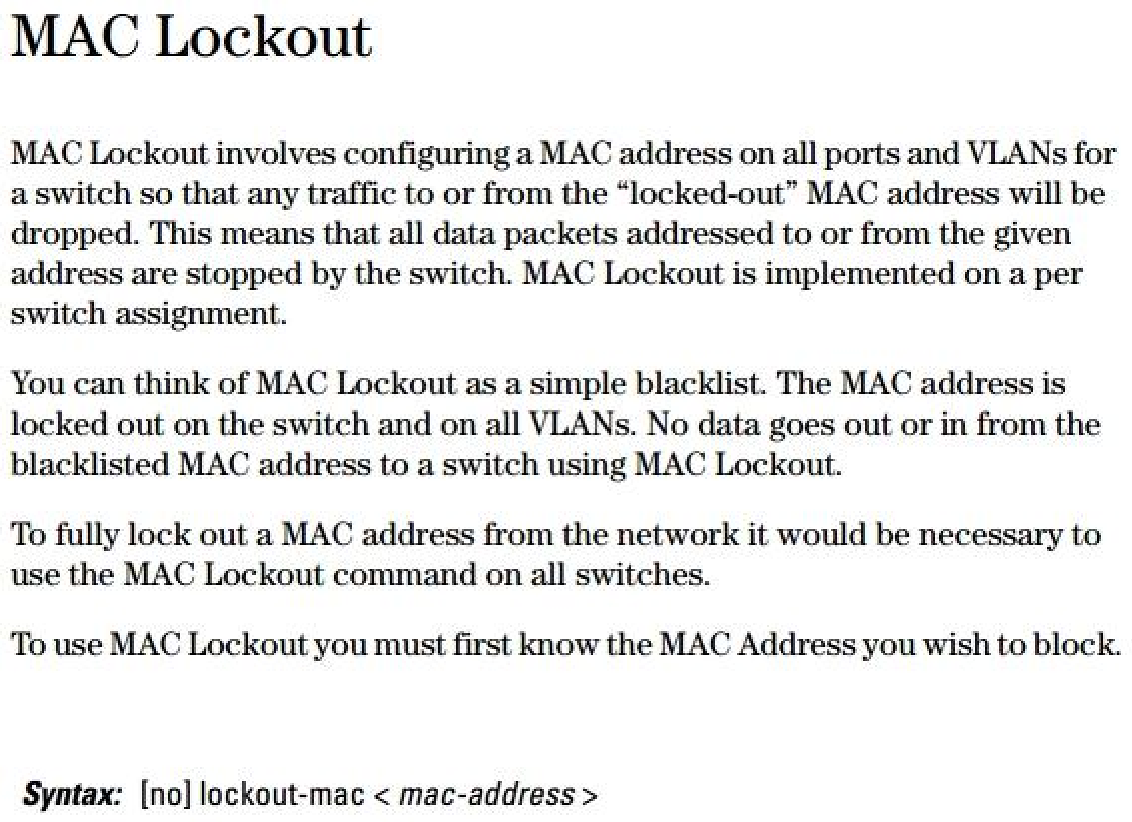

MAC Lockout

MAC Lockout involves configuring a MAC address on all ports and VLANs for a switch so that any traffic to or from the "locked-out" MAC address will be dropped. This means that all data packets addressed to or from the given address are stopped by the switch. MAC Lockout is implemented on a per switch assignment.

You can think of MAC Lockout as a simple blacklist. The MAC address is locked out on the switch and on all VLANs. No data goes out or in from the blacklisted MAC address to a switch using MAC Lockout.

To fully lock out a MAC address from the network it would be necessary to use the MAC Lockout command on all switches.

To use MAC Lockout you must first know the MAC Address you wish to block.

Syntax: [no] lockout-mac < *mac-address* >

QUESTION 54

View the exhibit. A network administrator has activated connection rate filtering on an HP 8200 zl Series switch with the throttle action and medium sensitivity. However, a server connected to port A1 in VLAN 10 (IP address 10.1.10.10) needs to establish many connections with other backend servers as part of its typical behavior. The ports that this server uses are TCP 50000-50020. Based on the information provided in the exhibit, how can the network administrator ensure that this traffic is never blocked while leaving the current protections in effect?

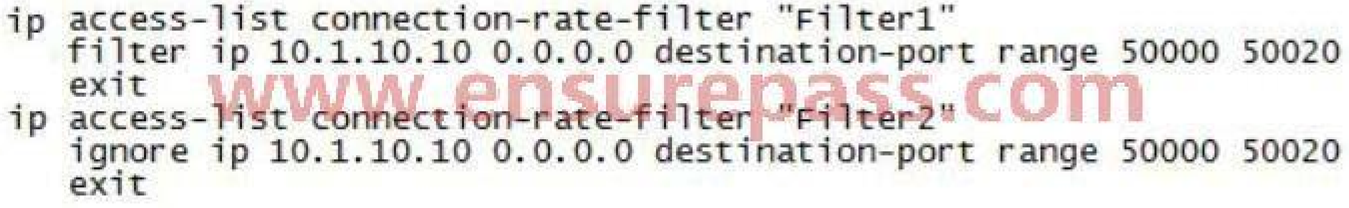

```
ip access-list connection-rate-filter "Filter1"
   filter ip 10.1.10.10 0.0.0.0 destination-port range 50000 50020
   exit
ip access-list connection-rate-filter "Filter2"
   ignore ip 10.1.10.10 0.0.0.0 destination-port range 50000 50020
   exit
```

A. Apply Filter2 to VLAN 10 as a connection rate filter.
B. Apply Filter1 to port A1 as a connection rate filter.
C. Apply Filter2 to port A1 as a connection rate filter.
D. Apply Filter1 to VLAN 10 as a connection rate filter.

Correct Answer: A
Section: QOS
Explanation

Explanation/Reference:
HP ProCurve Access Security Guide

Note: Command **connection-rate-filter** can not be applied to the physical interface of a switch! Only to VLAN-interface.

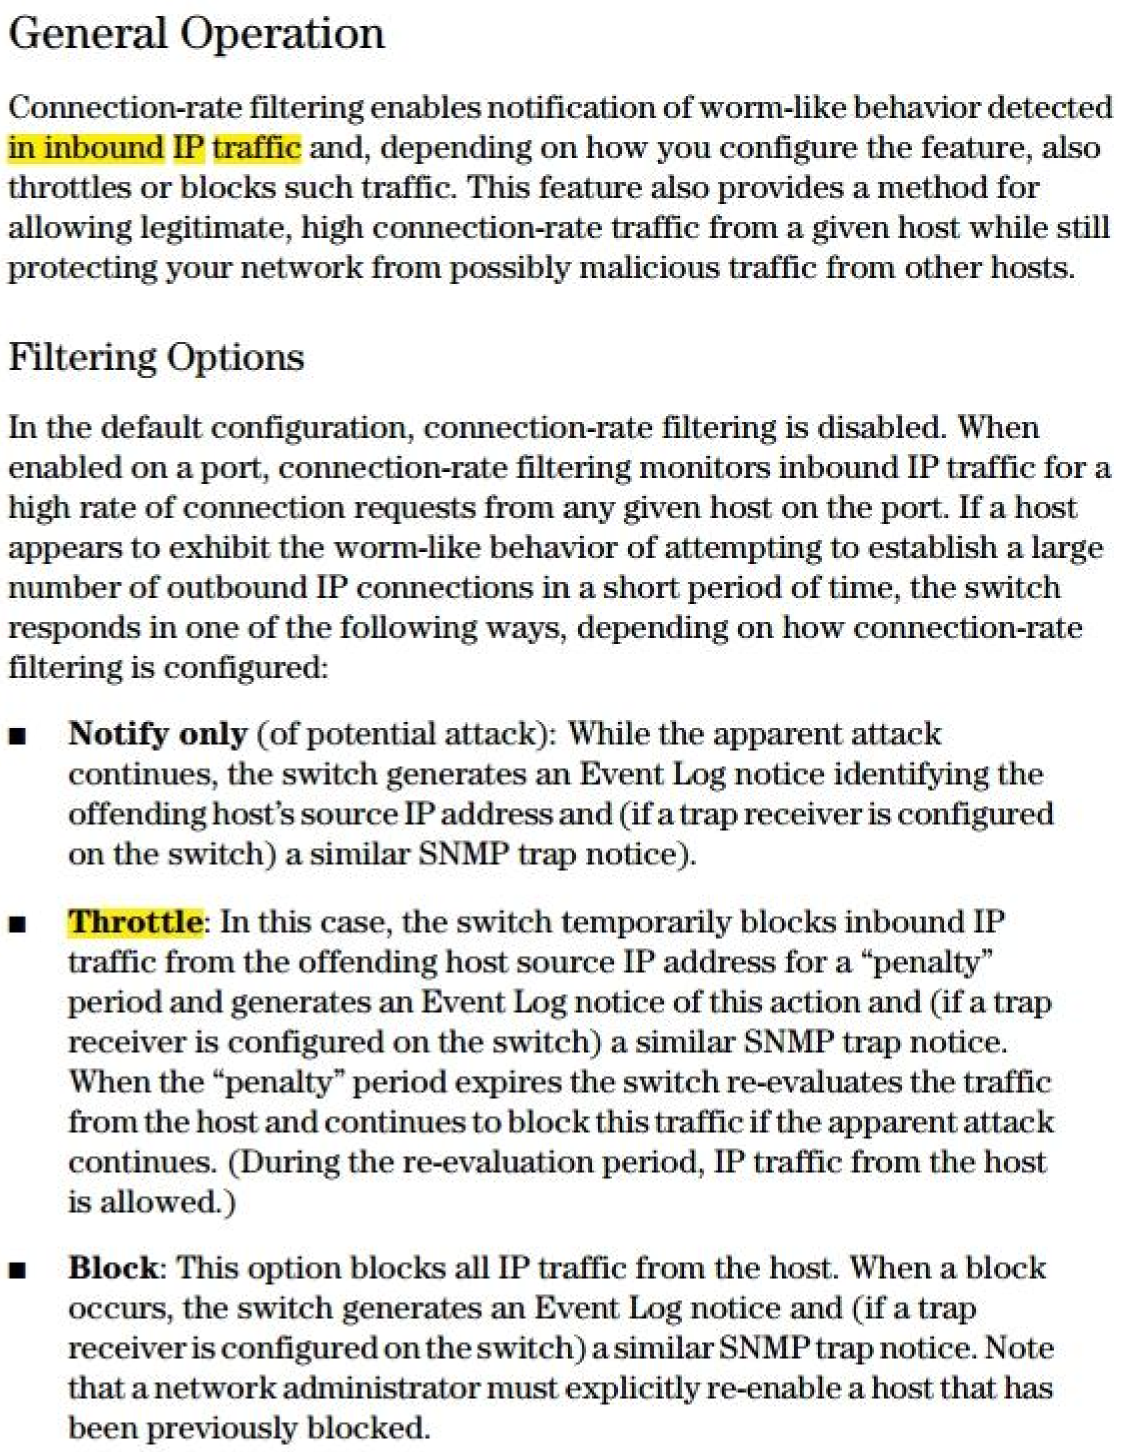

General Operation

Connection-rate filtering enables notification of worm-like behavior detected in inbound IP traffic and, depending on how you configure the feature, also throttles or blocks such traffic. This feature also provides a method for allowing legitimate, high connection-rate traffic from a given host while still protecting your network from possibly malicious traffic from other hosts.

Filtering Options

In the default configuration, connection-rate filtering is disabled. When enabled on a port, connection-rate filtering monitors inbound IP traffic for a high rate of connection requests from any given host on the port. If a host appears to exhibit the worm-like behavior of attempting to establish a large number of outbound IP connections in a short period of time, the switch responds in one of the following ways, depending on how connection-rate filtering is configured:

- **Notify only** (of potential attack): While the apparent attack continues, the switch generates an Event Log notice identifying the offending host's source IP address and (if a trap receiver is configured on the switch) a similar SNMP trap notice).

- **Throttle**: In this case, the switch temporarily blocks inbound IP traffic from the offending host source IP address for a "penalty" period and generates an Event Log notice of this action and (if a trap receiver is configured on the switch) a similar SNMP trap notice. When the "penalty" period expires the switch re-evaluates the traffic from the host and continues to block this traffic if the apparent attack continues. (During the re-evaluation period, IP traffic from the host is allowed.)

- **Block**: This option blocks all IP traffic from the host. When a block occurs, the switch generates an Event Log notice and (if a trap receiver is configured on the switch) a similar SNMP trap notice. Note that a network administrator must explicitly re-enable a host that has been previously blocked.

...

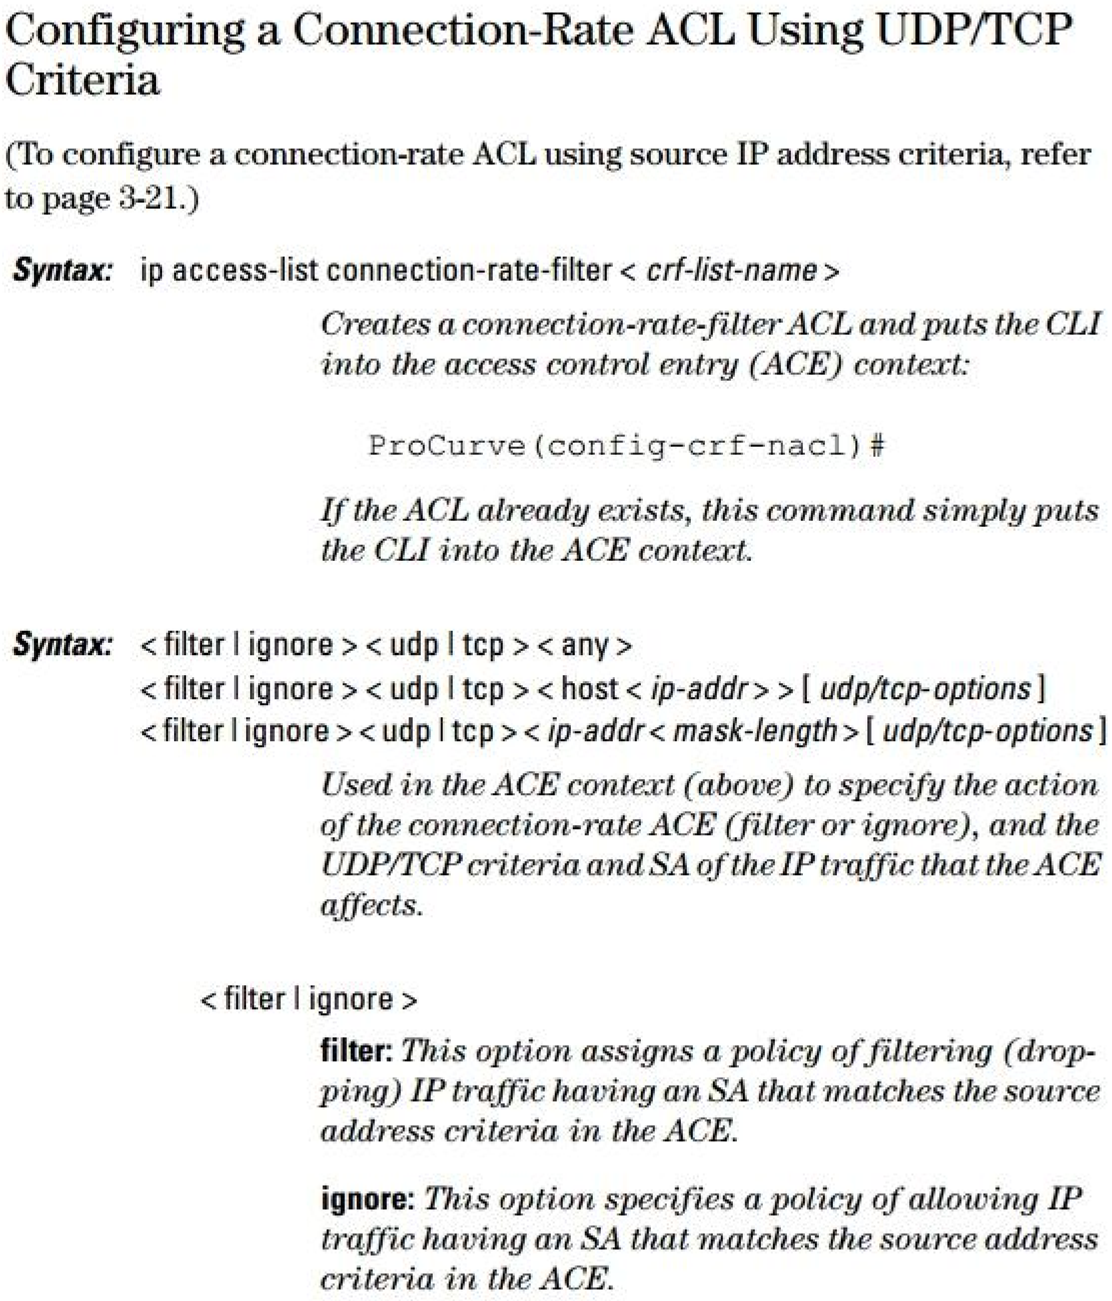

Configuring a Connection-Rate ACL Using UDP/TCP Criteria

(To configure a connection-rate ACL using source IP address criteria, refer to page 3-21.)

Syntax: ip access-list connection-rate-filter < *crf-list-name* >

Creates a connection-rate-filter ACL and puts the CLI into the access control entry (ACE) context:

```
ProCurve(config-crf-nacl)#
```

If the ACL already exists, this command simply puts the CLI into the ACE context.

Syntax: < filter | ignore > < udp | tcp > < any >
< filter | ignore > < udp | tcp > < host < *ip-addr* > > [*udp/tcp-options*]
< filter | ignore > < udp | tcp > < *ip-addr* < *mask-length* > [*udp/tcp-options*]

Used in the ACE context (above) to specify the action of the connection-rate ACE (filter or ignore), and the UDP/TCP criteria and SA of the IP traffic that the ACE affects.

< filter | ignore >

filter: *This option assigns a policy of filtering (dropping) IP traffic having an SA that matches the source address criteria in the ACE.*

ignore: *This option specifies a policy of allowing IP traffic having an SA that matches the source address criteria in the ACE.*

QUESTION 55
View the exhibit. Based on the configuration shown in the exhibit, does the Comware switch drop or permit each frame as it arrives on port Gigabit Ethernet 1/0/1 ? (Select three.)

A. DHCP offer are permitted.
B. DHCP offer are dropped.
C. DHCP Discovery (option 82 set) are permitted.
D. DHCP Discovery (option 82 set) are dropped.
E. ARP response for 10.1.10.2 is permitted.
F. ARP response for 10.1.10.2 is dropped.

Correct Answer: ACF
Section: Other
Explanation

Explanation/Reference:
Remember

QUESTION 56
Which security protocol introduces vulnerabilities because the password is sent in plaintext and can be intercepted and easily read?

A. WEP
B. EAP
C. PAP
D. CHAP

Correct Answer: C
Section: Security
Explanation

Explanation/Reference:
Password authentication protocol

A password authentication protocol (PAP) is an authentication protocol that uses a password.

PAP is used by Point to Point Protocol to validate users before allowing them access to server resources. Almost all network operating system remote servers support PAP.

PAP transmits unencrypted ASCII passwords over the network and is therefore considered unsecure. It is used as a last resort when the remote server does not support a stronger authentication protocol, like CHAP or EAP (the latter is actually a framework).

Password-based authentication is the protocol where two entities share a password in advance and use the password as the basis of authentication. Existing password authentication schemes can be categorized into two types: weak-password authentication schemes and strong-password authentication schemes. When compared to strong-password schemes, weak-password schemes tend to have lighter computational overhead, the designs are simpler, and implementation is easier, making them especially suitable for some constrained environments.

QUESTION 57
Which security protocol requires the servers to use digital certificates?

A. AH and ESP for HTTP traffic
B. Telnet for console sessions
C. SSL and TLS for HTTP traffic
D. SSH for console sessions

Correct Answer: C
Section: Security
Explanation

Explanation/Reference:

QUESTION 58
An HP Provision ASIC switch implements connection rate filtering. When an endpoint violates the connection rate policy, its traffic is filtered for less than 30 seconds. The company wants violator traffic to be filtered for about one minute. How should the network administrator complete this task?

A. Increase the global connection rate filtering sensitivity to high.
B. Increase the penalty period associated with the low sensitivity.
C. Increase the global penalty period for connection rate filtering.
D. Configure the connection rate filtering action on the ports from throttle to block mode.

Correct Answer: A
Section: Security
Explanation

Explanation/Reference:

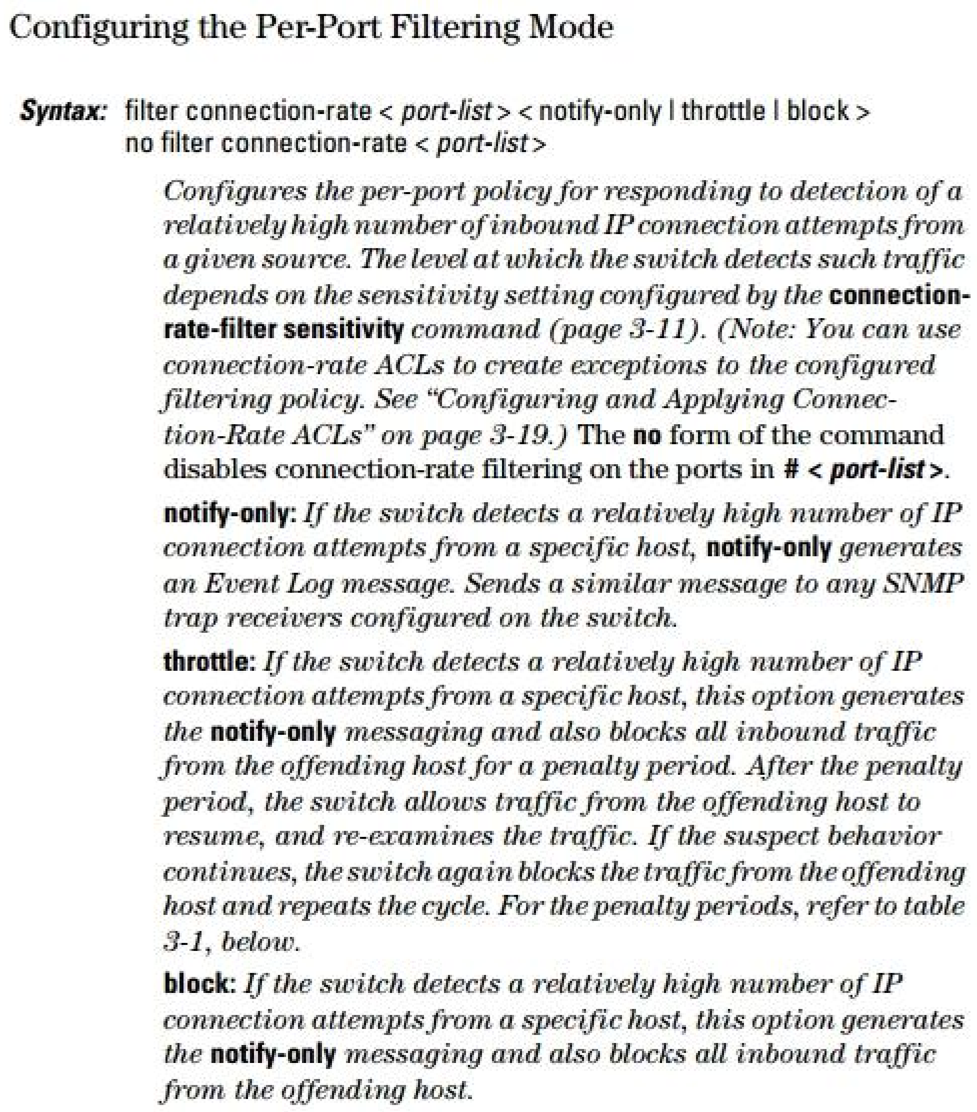

Configuring the Per-Port Filtering Mode

Syntax: filter connection-rate < *port-list* > < notify-only | throttle | block >
no filter connection-rate < *port-list* >

Configures the per-port policy for responding to detection of a relatively high number of inbound IP connection attempts from a given source. The level at which the switch detects such traffic depends on the sensitivity setting configured by the **connection-rate-filter sensitivity** *command (page 3-11). (Note: You can use connection-rate ACLs to create exceptions to the configured filtering policy. See "Configuring and Applying Connection-Rate ACLs" on page 3-19.)* The **no** form of the command disables connection-rate filtering on the ports in **# < *port-list* >**.

notify-only: *If the switch detects a relatively high number of IP connection attempts from a specific host,* **notify-only** *generates an Event Log message. Sends a similar message to any SNMP trap receivers configured on the switch.*

throttle: *If the switch detects a relatively high number of IP connection attempts from a specific host, this option generates the* **notify-only** *messaging and also blocks all inbound traffic from the offending host for a penalty period. After the penalty period, the switch allows traffic from the offending host to resume, and re-examines the traffic. If the suspect behavior continues, the switch again blocks the traffic from the offending host and repeats the cycle. For the penalty periods, refer to table 3-1, below.*

block: *If the switch detects a relatively high number of IP connection attempts from a specific host, this option generates the* **notify-only** *messaging and also blocks all inbound traffic from the offending host.*

QUESTION 59
Which technology should be used to tunnel multicast traffic securely across a network?

A. IPsec VPN

B. GRE over IPsec
C. GRE
D. MLD over IPsec for IPv6 or IGMP over IPsec for IPv4

Correct Answer: B
Section: Security
Explanation

Explanation/Reference:

QUESTION 60
A network administrator plans to use HP Intelligent Management Center (IMC) to manage a network with thousands of HP switches. The administrator wants to use IMC for functions such as configuring ACLs and VLANs on multiple switches at once. What must the administrator verify when configuring IMC to discover the switches?

A. that IMC is configured with the proper Telnet login username and password for the switches
B. that IMC is configured with the proper SNMP community or user credentials for read-write access to the switches
C. that IMC is configured with the proper SNMP trap credentials for communicating with the switches
D. that IMC is configured with the proper SSH login username and password for the switches

Correct Answer: B
Section: IMC
Explanation

Explanation/Reference:

QUESTION 61
A network administrator is deploying HP 5800 switches at the access layer and wants to manage them with HP Intelligent Management Center (IMC). When SNMP is enabled on these switches, which version of SNMP is implemented by default?

A. SNMP version 3
B. SNMP version 1
C. SNMP version 2c with backward compatibility with version 1
D. SNMP version 2c

Correct Answer: A
Section: IMC
Explanation

Explanation/Reference:

QUESTION 62
A network administrator wants to configure a Comware switch to send HP Intelligent Management Center (IMC) sFlow samples. Where does a Comware switch direct sFlow samples?

A. To the multicast address defined for sFlow collector
B. To the IP address configured as the sFlow agent
C. To the IP address configured as the sFlow trap receiver
D. To the IP address configured as the sFlow collector

Correct Answer: D
Section: IMC
Explanation

Explanation/Reference:

QUESTION 63
You are configuring IP multicast on a customer network that has all HP E5400 zl switches installed. All VLANs have been created, and IP addresses and OSPF have been assigned to all routed interfaces. IGMP has been enabled for all VLANs that will support multicast hosts.

To enable PIM, you issue the following commands:

```
E5406(config)# ip multicast-routing
E5406(config)# router pim
```

What is the remaining step you must take in this process?

A. Enable PIM for every VLAN that will support Layer 3 multicast.
B. Enable IGMP at the global configuration level.
C. Enable PIM only on VLAN interfaces that lead to other routers.
D. Enable sparse mode for every VLAN that will support Layer 3 multicast.

Correct Answer: A
Section: Multicast
Explanation

Explanation/Reference:

The necessary steps for PIM-DM:

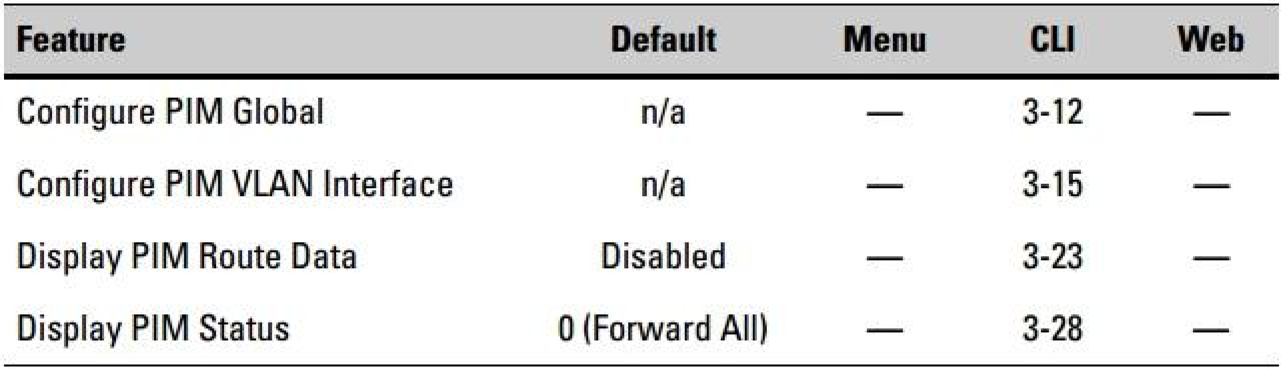

Feature	Default	Menu	CLI	Web
Configure PIM Global	n/a	—	3-12	—
Configure PIM VLAN Interface	n/a	—	3-15	—
Display PIM Route Data	Disabled	—	3-23	—
Display PIM Status	0 (Forward All)	—	3-28	—

The necessary steps for PIM-SM:

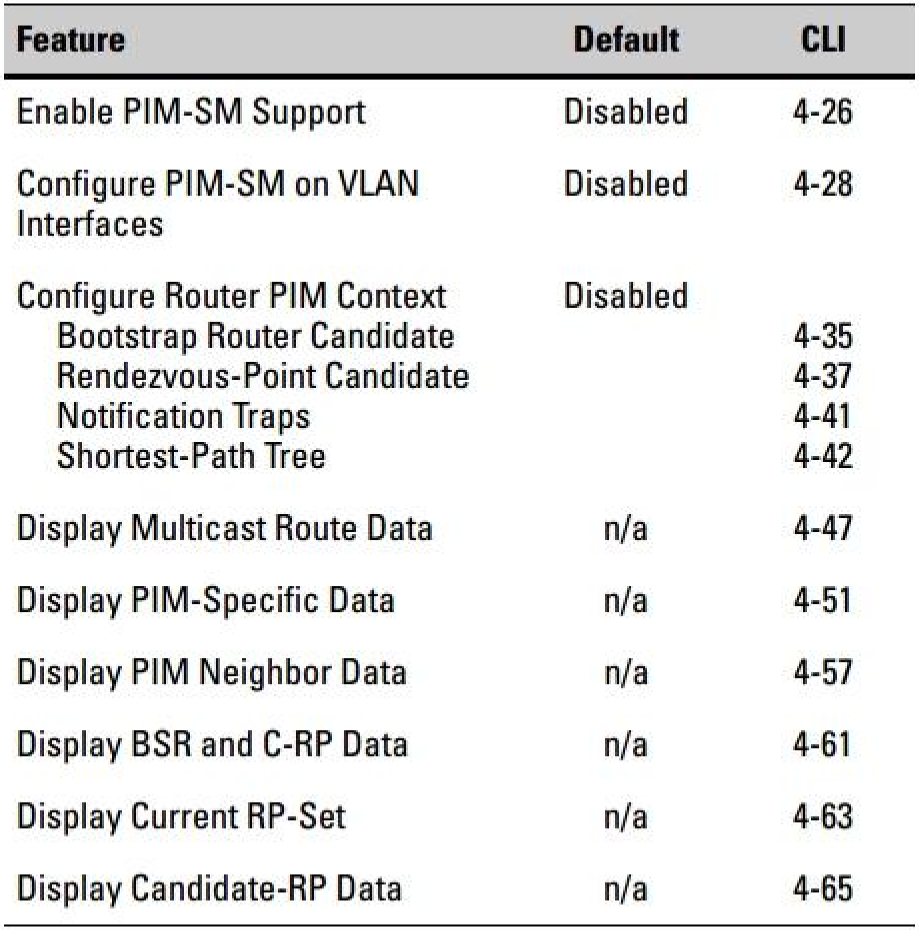

Feature	Default	CLI
Enable PIM-SM Support	Disabled	4-26
Configure PIM-SM on VLAN Interfaces	Disabled	4-28
Configure Router PIM Context	Disabled	
Bootstrap Router Candidate		4-35
Rendezvous-Point Candidate		4-37
Notification Traps		4-41
Shortest-Path Tree		4-42
Display Multicast Route Data	n/a	4-47
Display PIM-Specific Data	n/a	4-51
Display PIM Neighbor Data	n/a	4-57
Display BSR and C-RP Data	n/a	4-61
Display Current RP-Set	n/a	4-63
Display Candidate-RP Data	n/a	4-65

QUESTION 64
While onsite configuring an HP E5400 zl switch, you enter ipv6 enable in VLAN 1 configuration context. You type the command show ipv6 at the switch CLI, and receive the result shown in the exhibit. What is meant by the Address Status column?

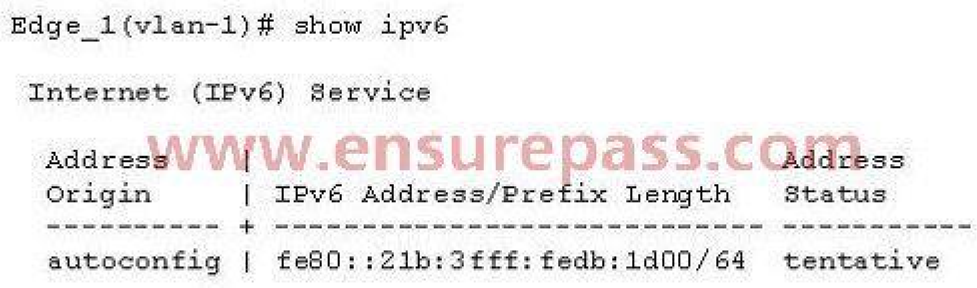

```
Edge_1(vlan-1)# show ipv6

 Internet (IPv6) Service

  Address    |                               Address
  Origin     | IPv6 Address/Prefix Length    Status
  ---------- + ---------------------------- -----------
  autoconfig | fe80::21b:3fff:fedb:1d00/64   tentative
```

A. The switch has secured its IPv6 address.
B. The switch is waiting for DAD to verify its IPv6 address.
C. The switch is waiting for an additional configuration command to be entered.
D. The switch is waiting for a router advertisement to verify its IPv6 address.

Correct Answer: B
Section: IPv6
Explanation

Explanation/Reference:
HP ProCurve Switch Software IPv6 Configuration Guide

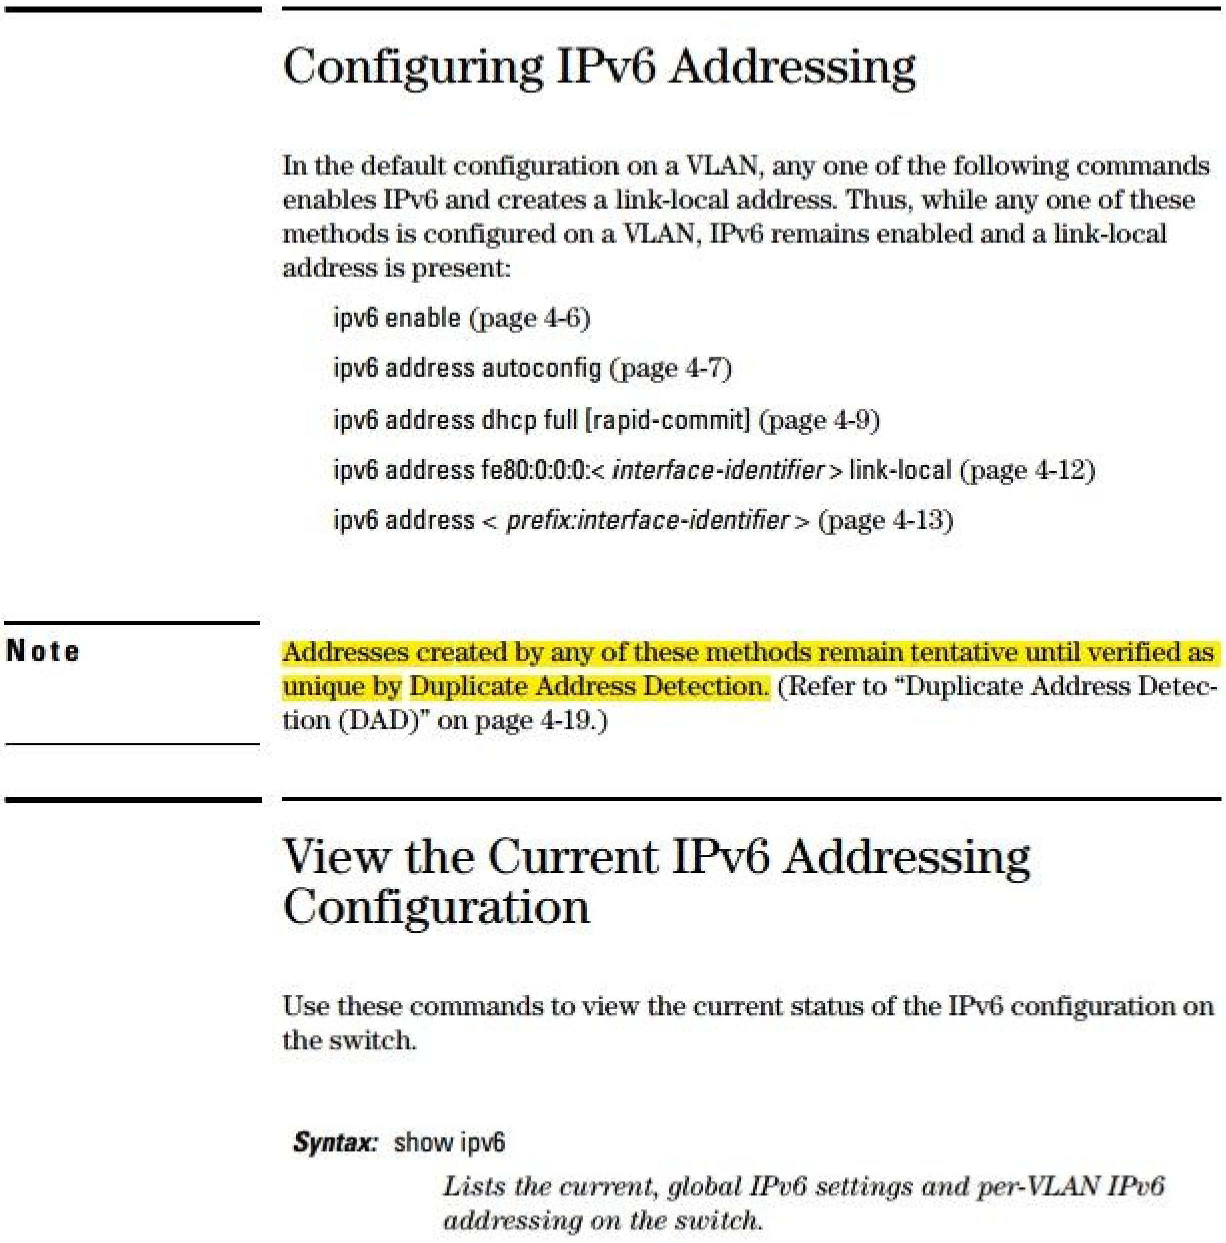

Configuring IPv6 Addressing

In the default configuration on a VLAN, any one of the following commands enables IPv6 and creates a link-local address. Thus, while any one of these methods is configured on a VLAN, IPv6 remains enabled and a link-local address is present:

ipv6 enable (page 4-6)

ipv6 address autoconfig (page 4-7)

ipv6 address dhcp full [rapid-commit] (page 4-9)

ipv6 address fe80:0:0:0:< *interface-identifier* > link-local (page 4-12)

ipv6 address < *prefix:interface-identifier* > (page 4-13)

Note

Addresses created by any of these methods remain tentative until verified as unique by Duplicate Address Detection. (Refer to "Duplicate Address Detection (DAD)" on page 4-19.)

View the Current IPv6 Addressing Configuration

Use these commands to view the current status of the IPv6 configuration on the switch.

Syntax: show ipv6

Lists the current, global IPv6 settings and per-VLAN IPv6 addressing on the switch.

...

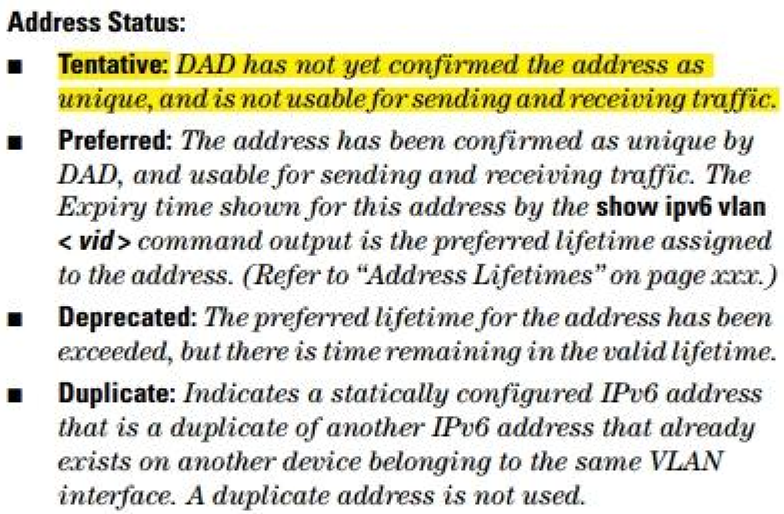

Address Status:

- **Tentative:** *DAD has not yet confirmed the address as unique, and is not usable for sending and receiving traffic.*
- **Preferred:** *The address has been confirmed as unique by DAD, and usable for sending and receiving traffic. The Expiry time shown for this address by the* **show ipv6 vlan < *vid* >** *command output is the preferred lifetime assigned to the address. (Refer to "Address Lifetimes" on page xxx.)*
- **Deprecated:** *The preferred lifetime for the address has been exceeded, but there is time remaining in the valid lifetime.*
- **Duplicate:** *Indicates a statically configured IPv6 address that is a duplicate of another IPv6 address that already exists on another device belonging to the same VLAN interface. A duplicate address is not used.*

QUESTION 65
An administrator of a network of A5800s has found evidence that an unauthorized device is gaining access to the network. All the administrator knows about the device is its MAC address. Assuming that the action she takes does not require special actions for guest devices, what can she do to keep the device from connecting to the network?

A. Implement MAC authentication
B. Lock out the MAC for the device, using the mac-address command
C. Lock out the MAC for the device, and statically assign the device to the 'null' interface using the static-mac command
D. Implement Port Security in the autoLearn mode

Correct Answer: B
Section: Security
Explanation

Explanation/Reference:

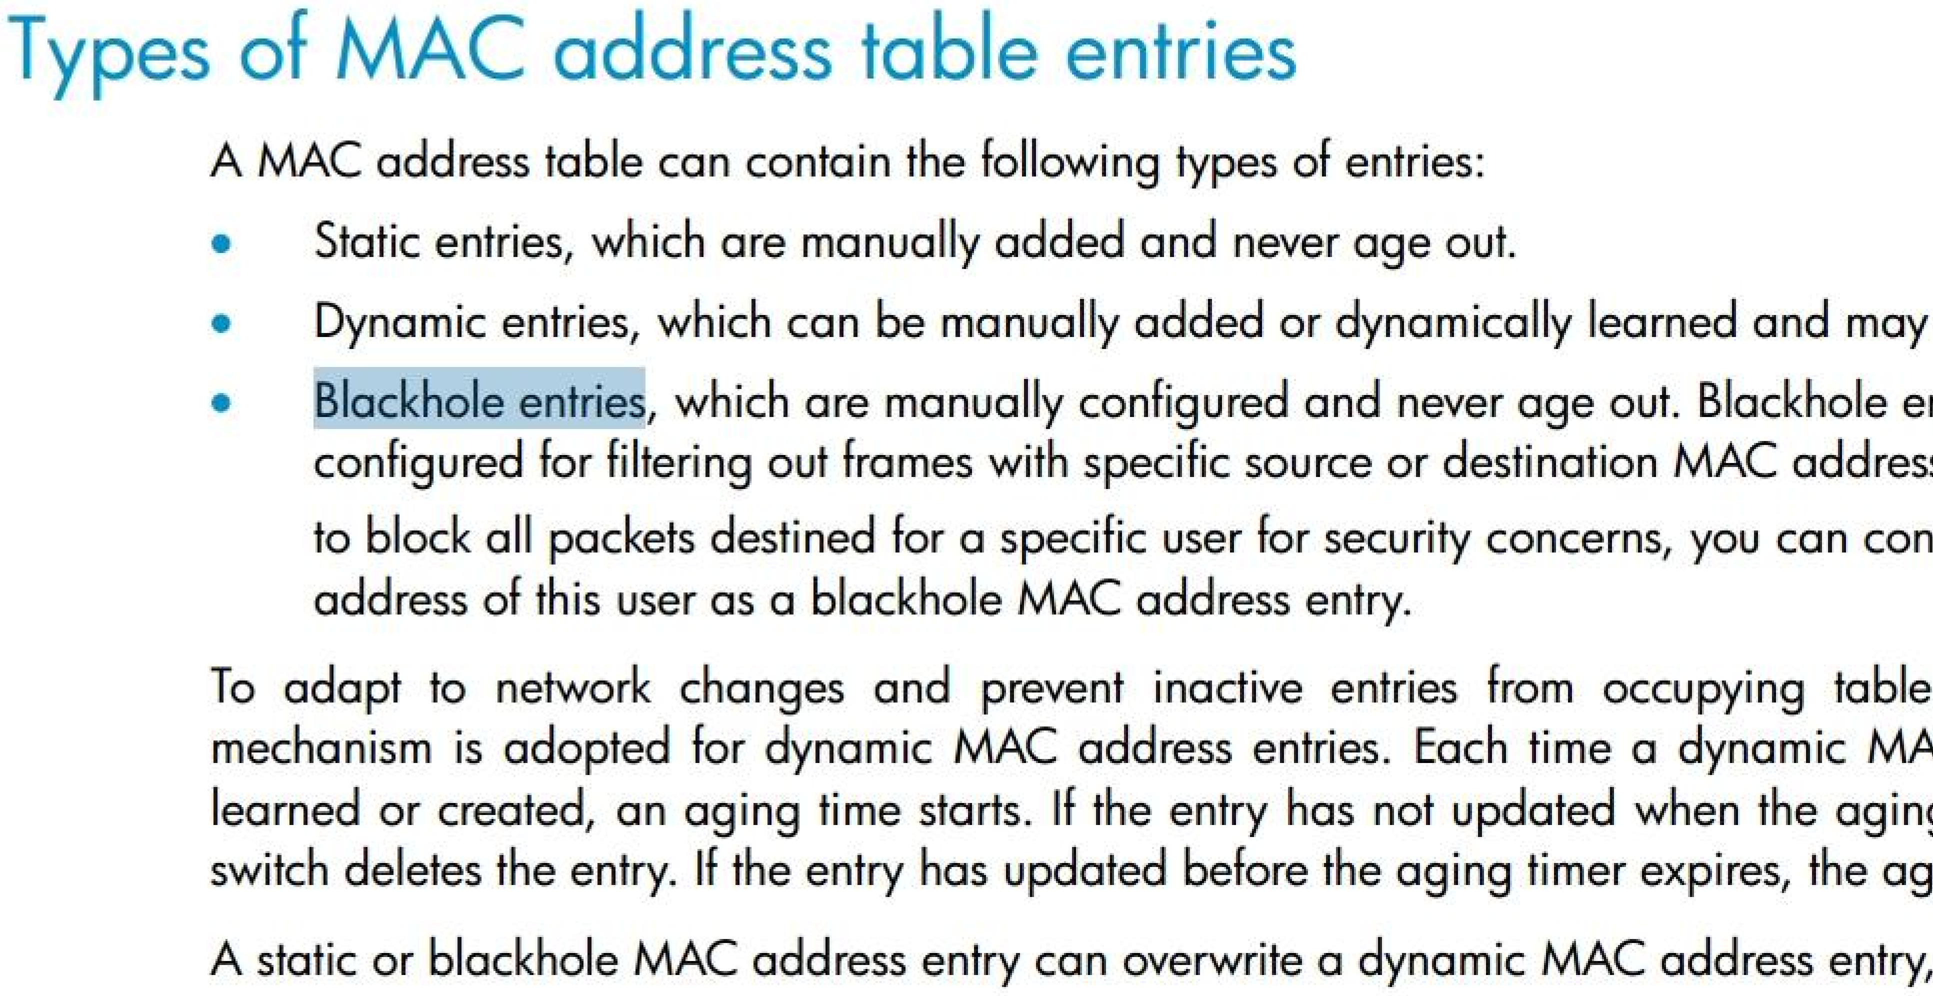

Types of MAC address table entries

A MAC address table can contain the following types of entries:

- Static entries, which are manually added and never age out.
- Dynamic entries, which can be manually added or dynamically learned and may
- Blackhole entries, which are manually configured and never age out. Blackhole e configured for filtering out frames with specific source or destination MAC addres

  to block all packets destined for a specific user for security concerns, you can con address of this user as a blackhole MAC address entry.

To adapt to network changes and prevent inactive entries from occupying table mechanism is adopted for dynamic MAC address entries. Each time a dynamic MA learned or created, an aging time starts. If the entry has not updated when the agin switch deletes the entry. If the entry has updated before the aging timer expires, the ag

A static or blackhole MAC address entry can overwrite a dynamic MAC address entry,

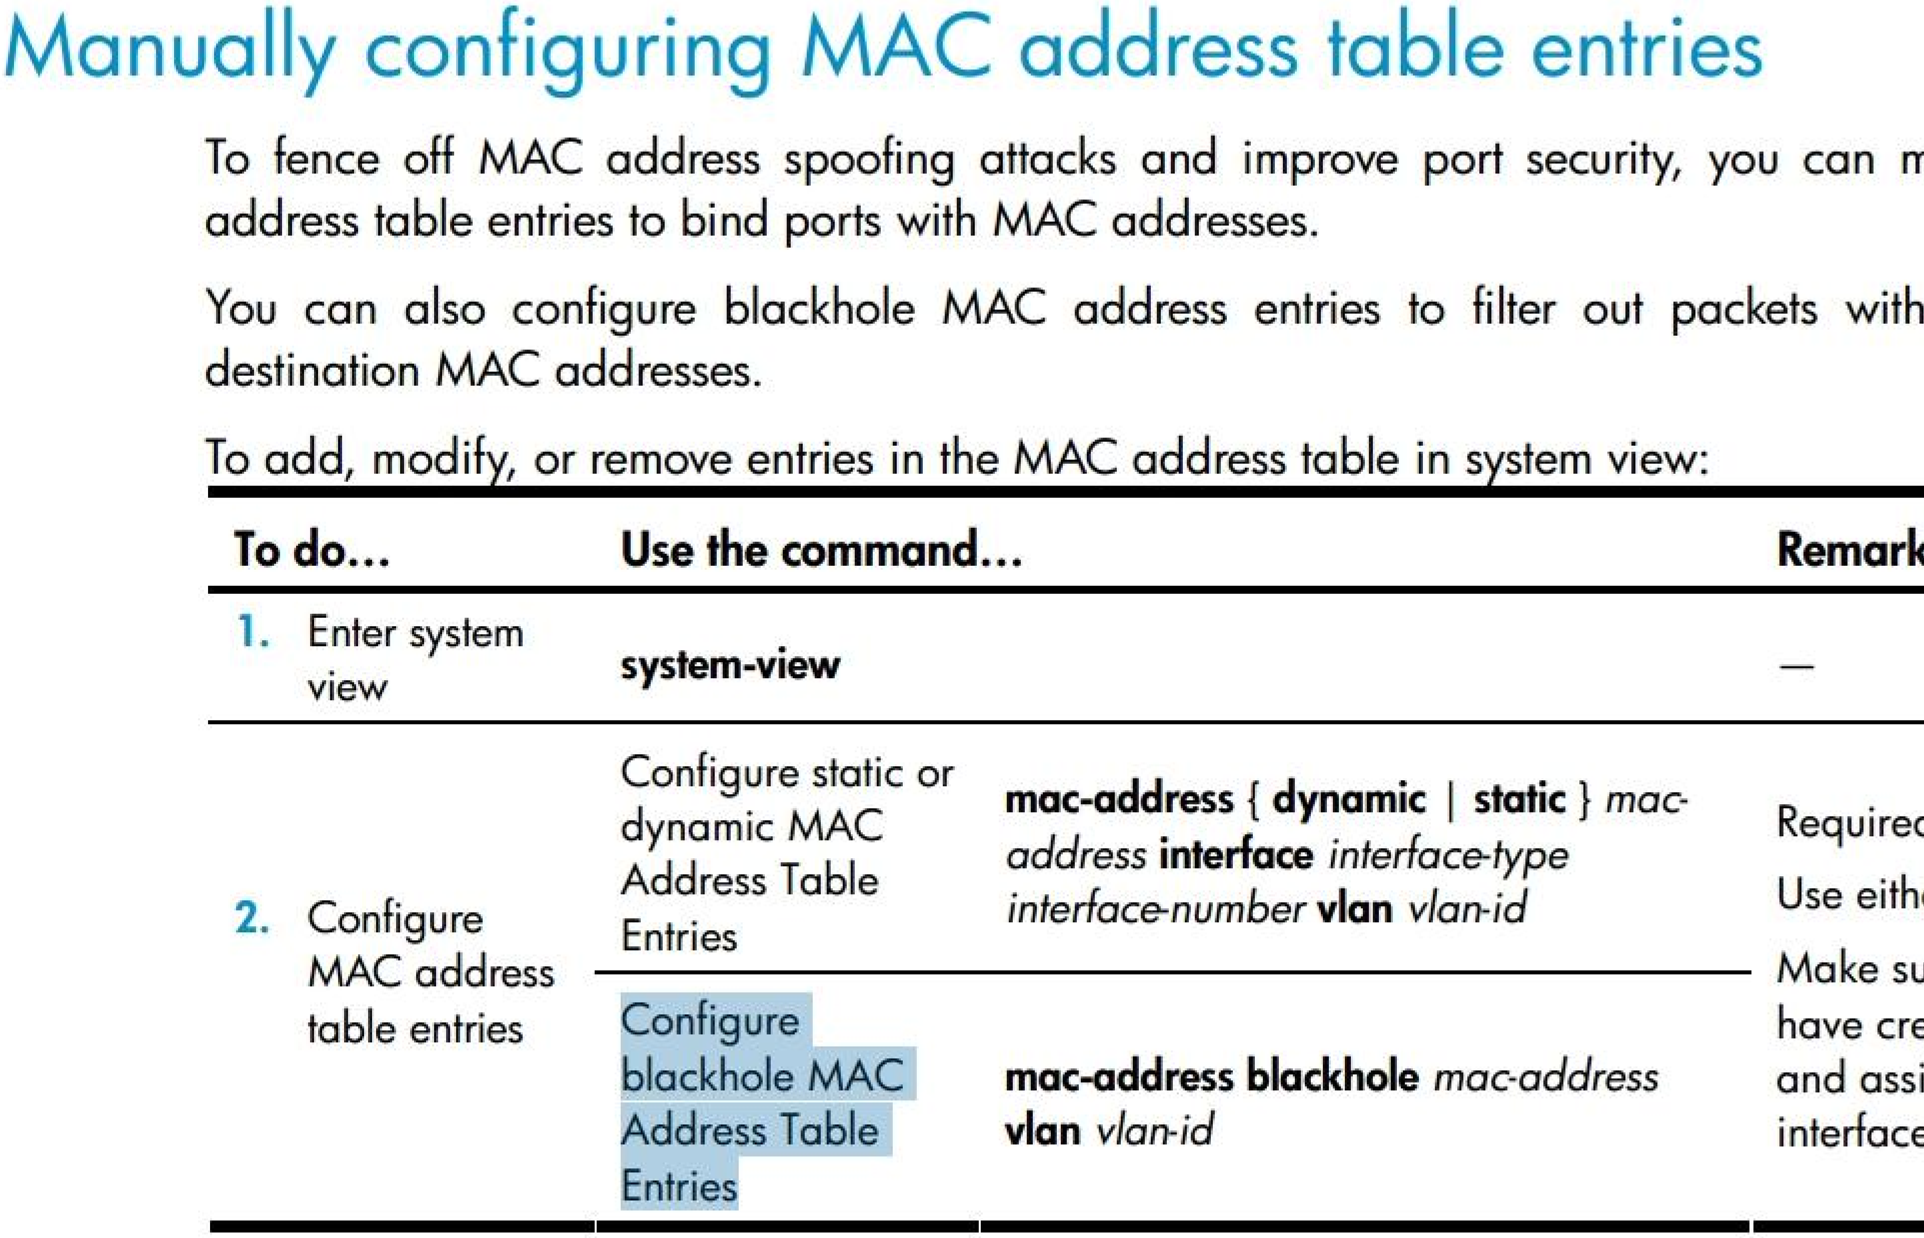

Manually configuring MAC address table entries

To fence off MAC address spoofing attacks and improve port security, you can m
address table entries to bind ports with MAC addresses.

You can also configure blackhole MAC address entries to filter out packets with
destination MAC addresses.

To add, modify, or remove entries in the MAC address table in system view:

To do...	Use the command...		Remark
1. Enter system view	**system-view**		—
2. Configure MAC address table entries	Configure static or dynamic MAC Address Table Entries	**mac-address** { **dynamic** | **static** } *mac-address* **interface** *interface-type interface-number* **vlan** *vlan-id*	Requirec Use eith Make su have cre and assi interface
	Configure blackhole MAC Address Table Entries	**mac-address blackhole** *mac-address* **vlan** *vlan-id*	

QUESTION 66
Which statement is true about the HP E5400 zl switch Connection-rate Filtering feature?

A. Any outbound traffic destined for a host that has been throttled or blocked is permitted.
B. When enabled, it is automatically globally activated.
C. It uses sFlow traffic sampling to determine whether traffic activity represents an intrusion.
D. It detects threats with both a signature-based engine and an anomaly-based engine that can detect zero day attacks.

Correct Answer: A
Section: Security
Explanation

Explanation/Reference:

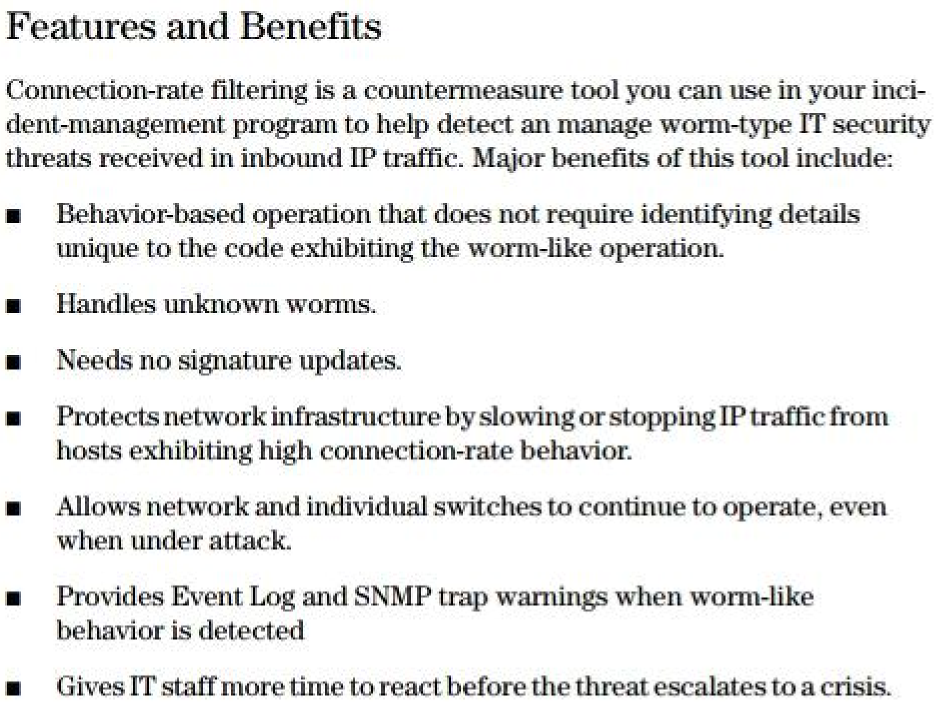

Features and Benefits

Connection-rate filtering is a countermeasure tool you can use in your incident-management program to help detect an manage worm-type IT security threats received in inbound IP traffic. Major benefits of this tool include:

- Behavior-based operation that does not require identifying details unique to the code exhibiting the worm-like operation.
- Handles unknown worms.
- Needs no signature updates.
- Protects network infrastructure by slowing or stopping IP traffic from hosts exhibiting high connection-rate behavior.
- Allows network and individual switches to continue to operate, even when under attack.
- Provides Event Log and SNMP trap warnings when worm-like behavior is detected
- Gives IT staff more time to react before the threat escalates to a crisis.

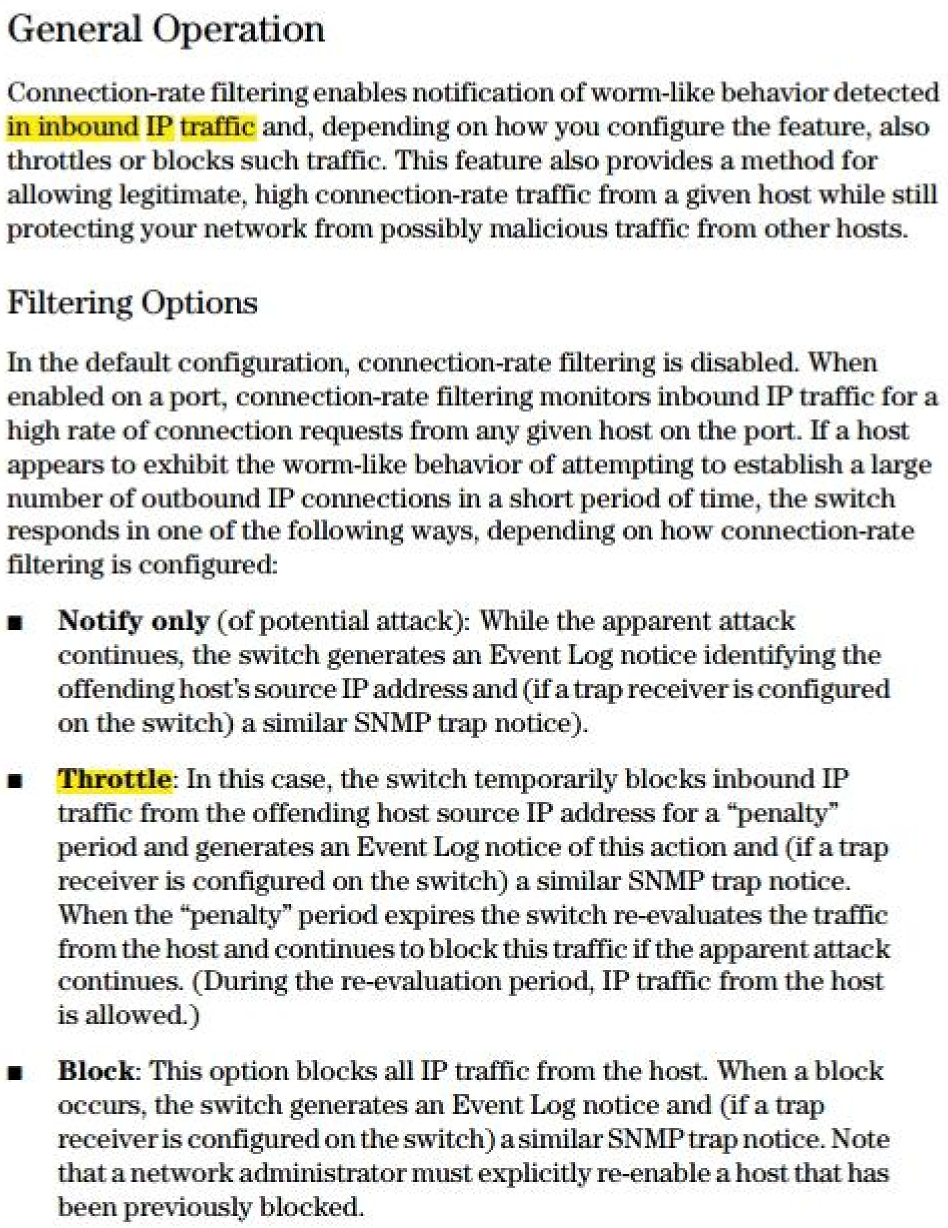

General Operation

Connection-rate filtering enables notification of worm-like behavior detected in inbound IP traffic and, depending on how you configure the feature, also throttles or blocks such traffic. This feature also provides a method for allowing legitimate, high connection-rate traffic from a given host while still protecting your network from possibly malicious traffic from other hosts.

Filtering Options

In the default configuration, connection-rate filtering is disabled. When enabled on a port, connection-rate filtering monitors inbound IP traffic for a high rate of connection requests from any given host on the port. If a host appears to exhibit the worm-like behavior of attempting to establish a large number of outbound IP connections in a short period of time, the switch responds in one of the following ways, depending on how connection-rate filtering is configured:

- **Notify only** (of potential attack): While the apparent attack continues, the switch generates an Event Log notice identifying the offending host's source IP address and (if a trap receiver is configured on the switch) a similar SNMP trap notice).
- **Throttle**: In this case, the switch temporarily blocks inbound IP traffic from the offending host source IP address for a "penalty" period and generates an Event Log notice of this action and (if a trap receiver is configured on the switch) a similar SNMP trap notice. When the "penalty" period expires the switch re-evaluates the traffic from the host and continues to block this traffic if the apparent attack continues. (During the re-evaluation period, IP traffic from the host is allowed.)
- **Block**: This option blocks all IP traffic from the host. When a block occurs, the switch generates an Event Log notice and (if a trap receiver is configured on the switch) a similar SNMP trap notice. Note that a network administrator must explicitly re-enable a host that has been previously blocked.

QUESTION 67

An administrator needs to create an ACL to block traffic from 192.168.10.23 to all destinations. Which type of ACL can be used to meet this need?

A. a standard ACL applied as Routed ACL (RACL)
B. an extended ACL applied as a Routed ACL (RACL)
C. a standard ACL applied as a VLAN ACL (VACL)
D. a standard ACL applied as a dynamic ACL

Correct Answer: C
Section: Security
Explanation

Explanation/Reference:

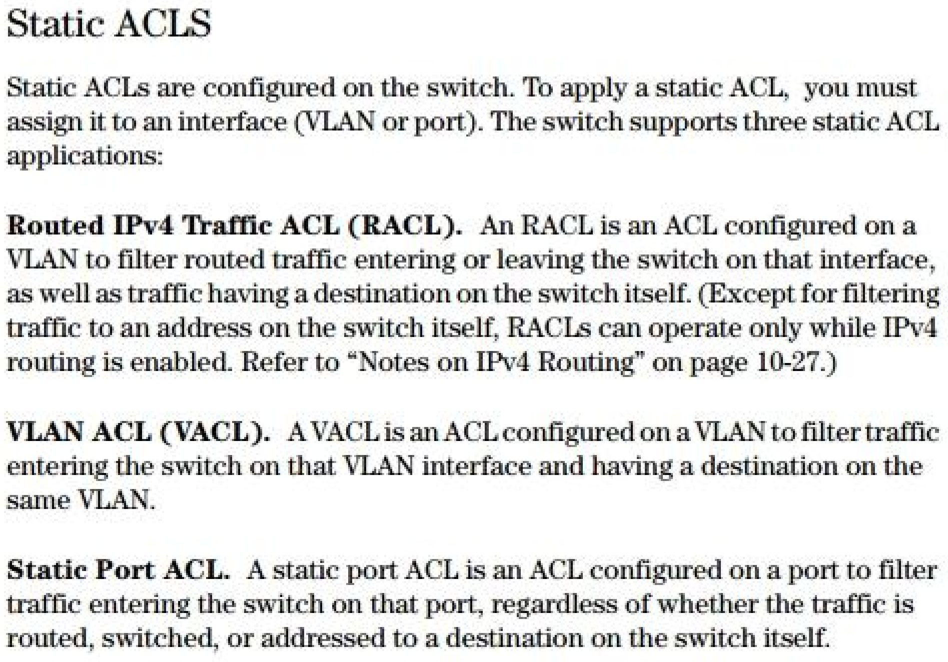

Static ACLS

Static ACLs are configured on the switch. To apply a static ACL, you must assign it to an interface (VLAN or port). The switch supports three static ACL applications:

Routed IPv4 Traffic ACL (RACL). An RACL is an ACL configured on a VLAN to filter routed traffic entering or leaving the switch on that interface, as well as traffic having a destination on the switch itself. (Except for filtering traffic to an address on the switch itself, RACLs can operate only while IPv4 routing is enabled. Refer to “Notes on IPv4 Routing” on page 10-27.)

VLAN ACL (VACL). A VACL is an ACL configured on a VLAN to filter traffic entering the switch on that VLAN interface and having a destination on the same VLAN.

Static Port ACL. A static port ACL is an ACL configured on a port to filter traffic entering the switch on that port, regardless of whether the traffic is routed, switched, or addressed to a destination on the switch itself.

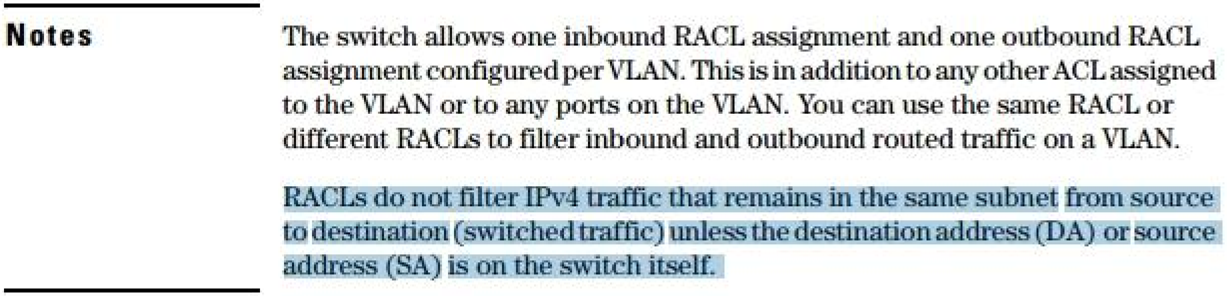

Notes

The switch allows one inbound RACL assignment and one outbound RACL assignment configured per VLAN. This is in addition to any other ACL assigned to the VLAN or to any ports on the VLAN. You can use the same RACL or different RACLs to filter inbound and outbound routed traffic on a VLAN.

RACLs do not filter IPv4 traffic that remains in the same subnet from source to destination (switched traffic) unless the destination address (DA) or source address (SA) is on the switch itself.

Standard ACL

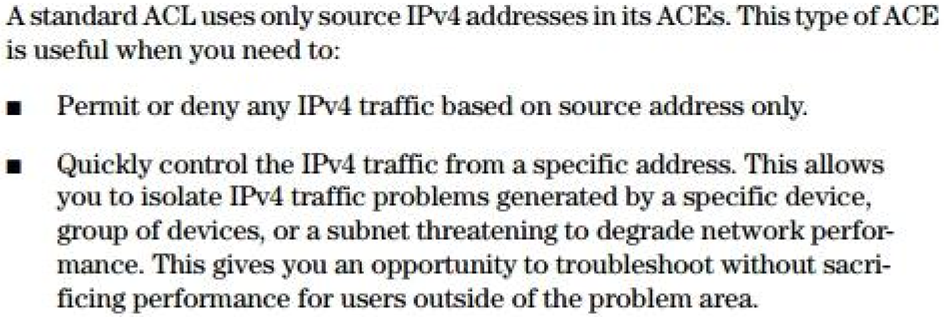

A standard ACL uses only source IPv4 addresses in its ACEs. This type of ACE is useful when you need to:

- Permit or deny any IPv4 traffic based on source address only.
- Quickly control the IPv4 traffic from a specific address. This allows you to isolate IPv4 traffic problems generated by a specific device, group of devices, or a subnet threatening to degrade network performance. This gives you an opportunity to troubleshoot without sacrificing performance for users outside of the problem area.

QUESTION 68
You are designing a network for a customer, and one of the requirements is to support multicast traffic. When defining configuration steps for the HP E5400 zl switches, when will both PIM and IGMP need to be configured?

A. when multicasts have multiple sources
B. when multicasts need to be routed
C. when multicasts have high usage among end users
D. when multicasts need to support large bandwidth

Correct Answer: B
Section: Multicast
Explanation

Explanation/Reference:

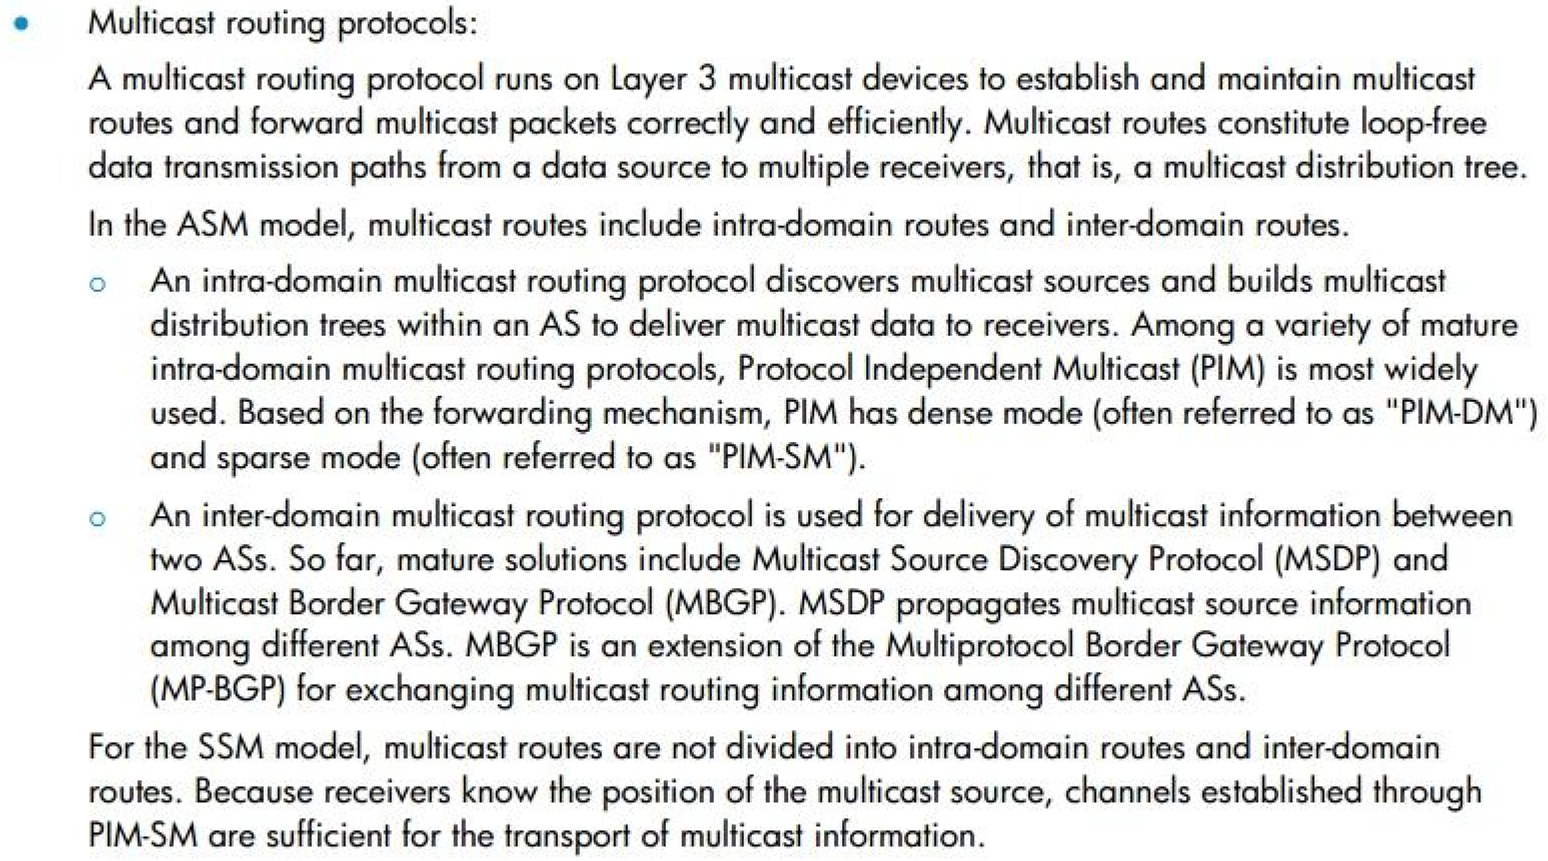

- Multicast routing protocols:

  A multicast routing protocol runs on Layer 3 multicast devices to establish and maintain multicast routes and forward multicast packets correctly and efficiently. Multicast routes constitute loop-free data transmission paths from a data source to multiple receivers, that is, a multicast distribution tree.

  In the ASM model, multicast routes include intra-domain routes and inter-domain routes.

  - An intra-domain multicast routing protocol discovers multicast sources and builds multicast distribution trees within an AS to deliver multicast data to receivers. Among a variety of mature intra-domain multicast routing protocols, Protocol Independent Multicast (PIM) is most widely used. Based on the forwarding mechanism, PIM has dense mode (often referred to as "PIM-DM") and sparse mode (often referred to as "PIM-SM").
  - An inter-domain multicast routing protocol is used for delivery of multicast information between two ASs. So far, mature solutions include Multicast Source Discovery Protocol (MSDP) and Multicast Border Gateway Protocol (MBGP). MSDP propagates multicast source information among different ASs. MBGP is an extension of the Multiprotocol Border Gateway Protocol (MP-BGP) for exchanging multicast routing information among different ASs.

  For the SSM model, multicast routes are not divided into intra-domain routes and inter-domain routes. Because receivers know the position of the multicast source, channels established through PIM-SM are sufficient for the transport of multicast information.

QUESTION 69
In A-Series switches, which VLAN types require the command local-proxy-arp enable for all the stations in the VLAN to be able to communicate to each other at Layer 3? (Select two.)

A. Port-based VLAN
B. SuperVLAN
C. Protocol-based VLAN
D. IP-subnet-based VLAN
E. MAC-address-based VLAN
F. Isolate-user VLAN

Correct Answer: BF
Section: VLANs
Explanation

Explanation/Reference:

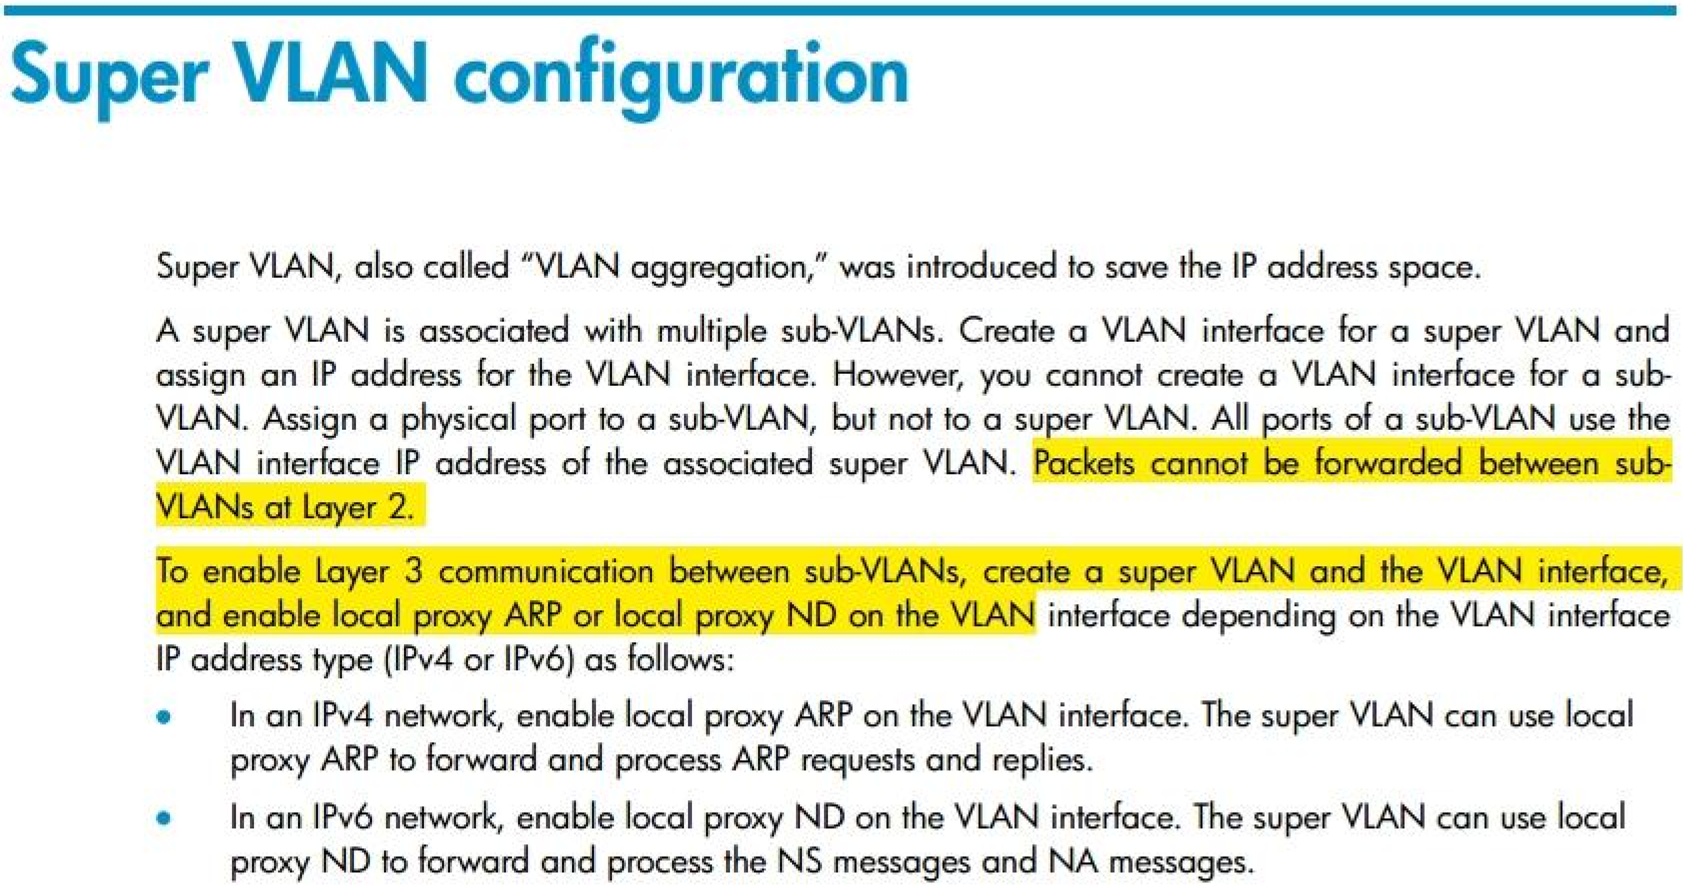

Super VLAN configuration

Super VLAN, also called "VLAN aggregation," was introduced to save the IP address space.

A super VLAN is associated with multiple sub-VLANs. Create a VLAN interface for a super VLAN and assign an IP address for the VLAN interface. However, you cannot create a VLAN interface for a sub-VLAN. Assign a physical port to a sub-VLAN, but not to a super VLAN. All ports of a sub-VLAN use the VLAN interface IP address of the associated super VLAN. Packets cannot be forwarded between sub-VLANs at Layer 2.

To enable Layer 3 communication between sub-VLANs, create a super VLAN and the VLAN interface, and enable local proxy ARP or local proxy ND on the VLAN interface depending on the VLAN interface IP address type (IPv4 or IPv6) as follows:

- In an IPv4 network, enable local proxy ARP on the VLAN interface. The super VLAN can use local proxy ARP to forward and process ARP requests and replies.
- In an IPv6 network, enable local proxy ND on the VLAN interface. The super VLAN can use local proxy ND to forward and process the NS messages and NA messages.

Isolate-user-VLAN configuration

An isolate-user-VLAN uses a two-tier VLAN structure. In this approach, both an isolate-user-VLAN and secondary VLANs are configured on the same device.

The isolate-user-VLAN implementation delivers the following benefits:

- Isolate-user-VLANs are mainly used for upstream data exchange. An isolate-user-VLAN can be associated with multiple secondary VLANs. Because the upstream device identifies only the isolate-user-VLAN and not the secondary VLANs, network configuration is simplified and VLAN resources are saved.
- You can isolate Layer 2 traffic from different users by assigning ports connected to them to different secondary VLANs. To enable communication between secondary VLANs associated with the same isolate-user-VLAN, you can enable local proxy ARP on the upstream device (such as Device A in Figure 42) to realize Layer 3 communication between the secondary VLANs.

As shown in Figure 42, the isolate-user-VLAN function is enabled on Device B. VLAN 10 is the isolate-user-VLAN. VLAN 2, VLAN 5, and VLAN 8 are secondary VLANs associated with VLAN 10 and are invisible to Device A.

Figure 42 An isolate-user-VLAN example

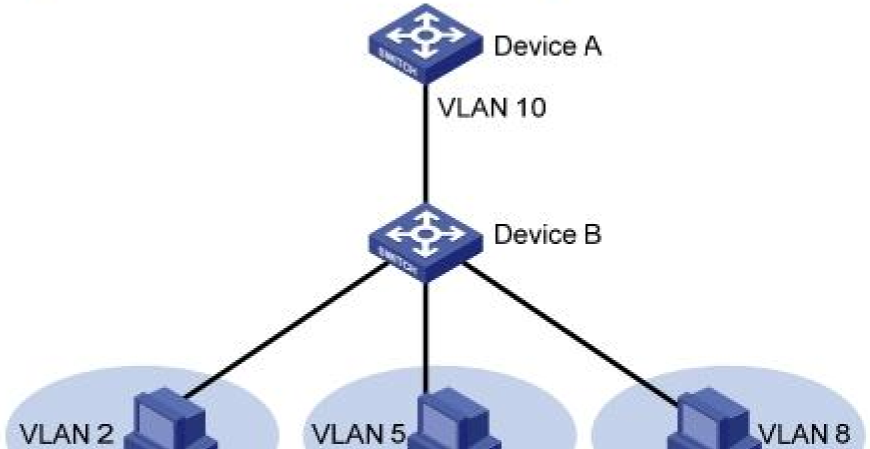
Device A
VLAN 10
Device B
VLAN 2
VLAN 5
VLAN 8