

# HP

## Exam HP0-Y47

### Deploying HP FlexNetwork Core Technologies

Version: 7.0

[ Total Questions: 55 ]

**Question No : 1**

A network administrator is completing an In-Service Software Upgrade (ISSU) for an Intelligent Resilient Framework (IRF) virtual devices. The device has two members. Each member has one management module. Member1 is currently the master. The administrator has initiated a rollback time for this upgrade.

When should the administrator accept the upgrade?

- A. After upgrading both members of the IRF virtual device
- B. After checking the new software's ISSU compatibility but before upgrading either member
- C. After upgrading member 1 but before switching over to and upgrading member 2
- D. After upgrading and switching over to member 2 but before upgrading member 1

**Answer: A**

**Explanation:** Performing an ISSU for a multi-member IRF fabric

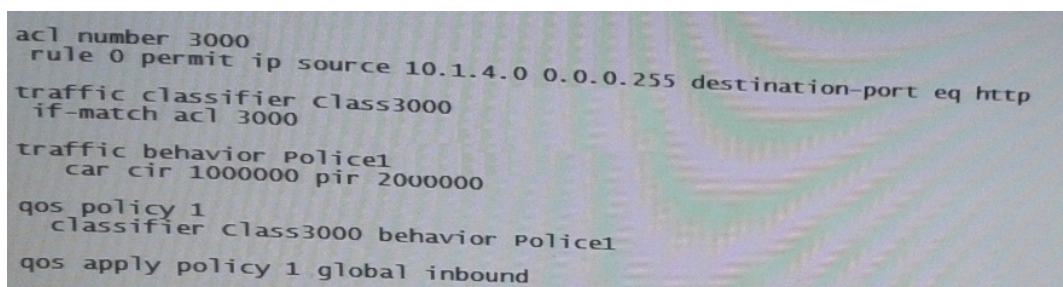
ISSU for a multi-member IRF fabric should be performed as per member in two steps: first upgrade a subordinate member, and then upgrade the master and the other subordinate members.

08-ISSU

configuration[http://www.h3c.com/portal/Technical\\_Support\\_\\_\\_Documents/Technical\\_Documents/Switches/H3C\\_S12500\\_Series\\_Switches/Configuration/Operation\\_Manual/H3C\\_S12500\\_CG-Release7128-6W710/01/201301/772597\\_1285\\_0.htm](http://www.h3c.com/portal/Technical_Support___Documents/Technical_Documents/Switches/H3C_S12500_Series_Switches/Configuration/Operation_Manual/H3C_S12500_CG-Release7128-6W710/01/201301/772597_1285_0.htm)

**Question No : 2**

Refer to the exhibit.



```
acl number 3000
 rule 0 permit ip source 10.1.4.0 0.0.0.255 destination-port eq http
traffic classifier class3000
 if-match acl 3000
traffic behavior policel
 car cir 1000000 pir 2000000
qos policy 1
 classifier class3000 behavior policel
qos apply policy 1 global inbound
```

This HP 10500 Switch Series is receiving an average of 1 Gbps of HTTP traffic from 10.1.4.0/24. The switch starts to receive an additional 1 Gbps of HTTP traffic from 10.1.4.0/24. How does the switch handle the traffic?

- A. It drops the traffic
- B. It forwards the traffic but marks it yellow (for a higher drop precedence)
- C. It forwards the traffic without remarking it in any way
- D. It forwards the traffic but marks it for forwarding in a lower priority queue

**Answer: C**

**Explanation:**

**Parameters**

**cir***committed-information-rate*: Specifies the committed information rate (CIR) in kbps.

**cbs***committed-burst-size*: Specifies the committed burst size (CBS) in bytes. The *committed-burst-size* argument ranges from 4000 to 16000000, the default is 4000.

**eb***excess-burst-size*: Specifies excess burst size (EBS) in bytes. The *excess-burst-size* argument ranges from 0 to 16000000, the default is 4000.

**pir***peak-information-rate*: Specifies the peak information rate (PIR) in kbps.

**green***action*: Specifies the action to be conducted for the traffic conforming to CIR. The *action* argument can be:

**discard**: Drops the packets.

**pass**: Forwards the packets.

**remark-dscp-pass***new-dscp*: Marks the packets with a new DSCP precedence and forwards them to their destinations. The *new-dscp* argument is in the range 0 to 63. By default, packets conforming to CIR are forwarded.

**red***action*: Specifies the action to be conducted for the traffic conforms to neither CIR nor PIR. The *action* argument can be:

**discard**: Drops the packets.

**pass**: Forwards the packets.

**remark-dscp-pass***new-dscp*: Marks the packets with a new DSCP precedence and forwards them to their destinations. The *new-dscp* argument is in the range 0 to 63. By default, packets conforming to neither CIR nor PIR are dropped.

**yellow***action*: Specifies the action to be conducted for the traffic conforms to PIR but does not conform to CIR. The *action* argument can be:

**discard**: Drops the packets.

**pass**: Forwards the packets.

**remark-dscp-pass***new-dscp*: Marks the packets with a new DSCP precedence and forwards them to their destinations. The *new-dscp* argument is in the range 0 to 63. By default, packets conforming to PIR but not conforming to CIR are forwarded.

**Question No : 3**

The network administrator is configuring Unidirectional Link Detection (UDLD) on trk1 on an HP Provision switch. The trk1 link aggregation group has these settings:

Interfaces = 48 and 49

Untagged VLAN = 1

Tagged VLANs = 2-4

How should the administrator complete the setup?

- A. Enable link-keepalive on both interfaces in the trunk (48 and 49).
- B. Enable link-keepalive on interface 48, specifying the VLAN as 1.
- C. On both interfaces in the trunk (48 and 49), enable link-keepalive, specifying the VLAN as 1.
- D. Enable link-keepalive on trk1

**Answer: C**

**Explanation:** \* The following commands allow you to configure UDLD via the CLI.

Syntax: [no] interface <port-list> link-keepalive

Enables UDLD on a port or range of ports

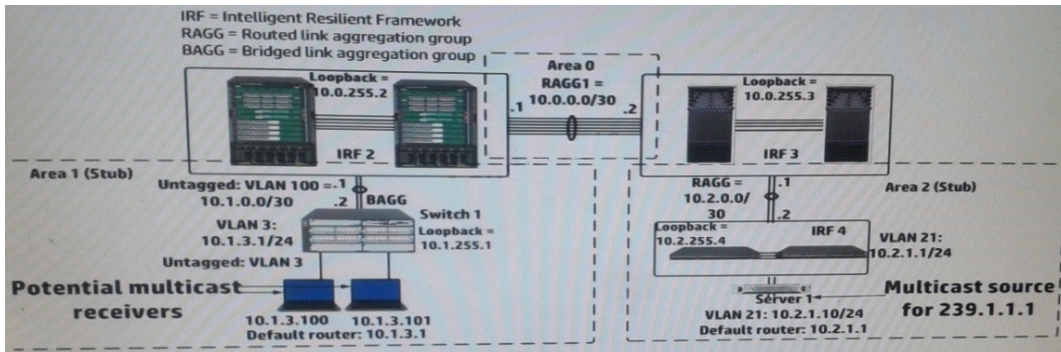
\* Ports enabled for UDLD exchange health-check packets once every five seconds (the link-keepalive interval). If a port does not receive a health-check packet from the port at the other end of the link within the keepalive interval, the port waits for four more intervals. If the port still does not receive a healthcheck packet after waiting for five intervals, the port concludes that the link has failed and blocks the UDLD-enabled port.

Reference:Port Status and Basic Configuration

[ftp://ftp.hp.com/pub/networking/software/2900yl-MCG-0207-T\\_12\\_XX-Chap10-PrtStat-BasicCfg.pdf](ftp://ftp.hp.com/pub/networking/software/2900yl-MCG-0207-T_12_XX-Chap10-PrtStat-BasicCfg.pdf)

**Question No : 4**

Refer to the exhibit.



A network administrator has begun configuring a Protocol Independent Multicast Sparse Mode (PIM-SM) solution on all of the network infrastructure devices shown in the exhibit. The administrator has selected the static method for configuring rendezvous (RPs) and wants to use 4 as RP.

Where must the administrator configure the static RP setting?

- A. On IRF 3 and IRF 4 only
- B. On switch 1, IRF 2, IRF 3, and IRF 4
- C. On IRF 2, IRF 3, and IRF 4 only
- D. On IRF 2 and IRF 3 only

**Answer: C**

**Explanation:**

In Static-RP configuration, you need to configure “*ip pim rp-address x.x.x.x*” command on every multicast enable router in your network including RP itself.

Reference:PIM-SM Static RP Configurations

<http://mrncciew.com/2013/01/19/pim-sm-static-rp-configurations/>

**Question No : 5**

A company wants to enforce source-specific multicasting (SSM) for several multicast

streams.

This is the configuration:

- ✍ Some multicast receivers support IGMPv3.
- ✍ Other receivers that need to stream support IGMPv2 only.
- ✍ The multicast receivers' default routers enable IGMPv3 on the interfaces associated with their subnets.
- ✍ The routers are correctly configured to enable the IGMPv3 receivers to receive the stream from the correct source

Which setting do the receivers' routers require to support the IGMPV2 devices?

- A. An SSM policy that selects the IGMPv2 devices and binds them to an RP address
- B. An SSM map that sets a rendezvous point (RP) for each multicast address
- C. An SSM policy that selects the appropriate multicast destinations and binds them to an RP address
- D. An SSM map that sets the correct source for each multicast address

**Answer: D**

**Explanation:** IGMP SSM mapping

The IGMP SSM mapping feature provides SSM support for receiver hosts that are running IGMPv1 or IGMPv2. This feature is implemented by configuring static IGMP SSM mappings on the IGMP-enabled routers.

The SSM model assumes that the IGMP-enabled routers have identified the desired multicast sources when receivers join multicast groups.

- A host running IGMPv3 can explicitly specify multicast source addresses in its reports.
- A host running IGMPv1 or IGMPv2, however, cannot specify multicast source addresses in its reports. In this case, you must configure the IGMP SSM mapping feature to translate the (\*, G) information in the IGMPv1 or IGMPv2 reports into (G, INCLUDE, (S1, S2...)) information

Reference:HP MSR Router Series, IP Multicast Configuration Guide(V7)

[http://h20565.www2.hp.com/hpsc/doc/public/display?sp4ts.oid=6796027&docId=emr\\_na-c04412136&docLocale=en\\_US](http://h20565.www2.hp.com/hpsc/doc/public/display?sp4ts.oid=6796027&docId=emr_na-c04412136&docLocale=en_US)

**Question No : 6 HOTSPOT**

Refer to the exhibit.

HP HP0-Y47 : Practice Test

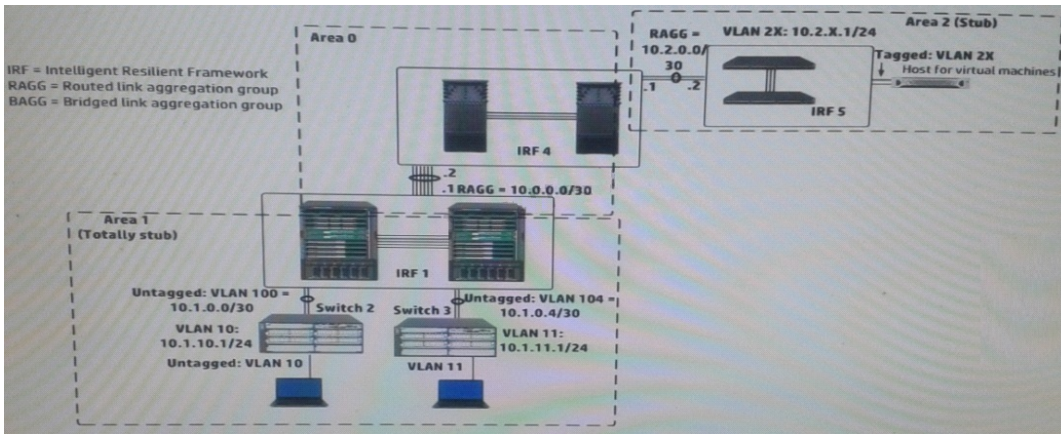


Exhibit 2

```
[IRF-1-ospf-1]display this
ospf 1
 area 0.0.0.0
  abr-summary 10.0.0.0 255.255.0.0 cost 1
  network 10.0.0.0 0.0.255.255
 area 0.0.0.1
  abr-summary 10.1.0.0 255.255.0.0 cost 1
  network 10.1.0.0 0.0.255.255
 stub no-summary
```

In exhibit 1, all infrastructure devices are implementing OSPF on the interfaces. Exhibit 1 also shows settings for OSPF areas. Exhibit 2 shows some additional OSPF settings IRF 1.

The network administrator enters this command on IRF 4:

```
[ IRF4-ospf-1-area-0.0.0.2 ] abr-summary 10.2.0.0 16
```

The administrator verifies that the solution is functioning correctly. Indicate whether the link state database (LSDB) on a device should include an LSA for 10.2.0.0/16.

**IRF 4 area 0 LSDB**

	▼
<b>Include LSA</b>	
<b>Does not include LSA</b>	

**Switch 2 area 1 LSDB**

	▼
<b>Include LSA</b>	
<b>Does not include LSA</b>	

Answer:

**IRF 4 area 0 LSDB**

	▼
<b>Include LSA</b>	
<b>Does not include LSA</b>	


**Switch 2 area 1 LSDB**

	▼
<b>Include LSA</b>	
<b>Does not include LSA</b>	

**Question No : 7**

Refer to the exhibit.





Switch 1 and Switch 2 have established an Interswitch Connection (ISC) and two distributed trunks. All of the switches have Multiple Spanning Tree Protocol (MSTP) enabled. Where does Switch 1 transmit or process Bridge Protocol Data Units (BPDUs)?

- A. On trk2 and the peer keepalive
- B. On trk1 and trk3
- C. On trk1, trk2, trk3
- D. On trk2

**Answer: C**

**Explanation:**

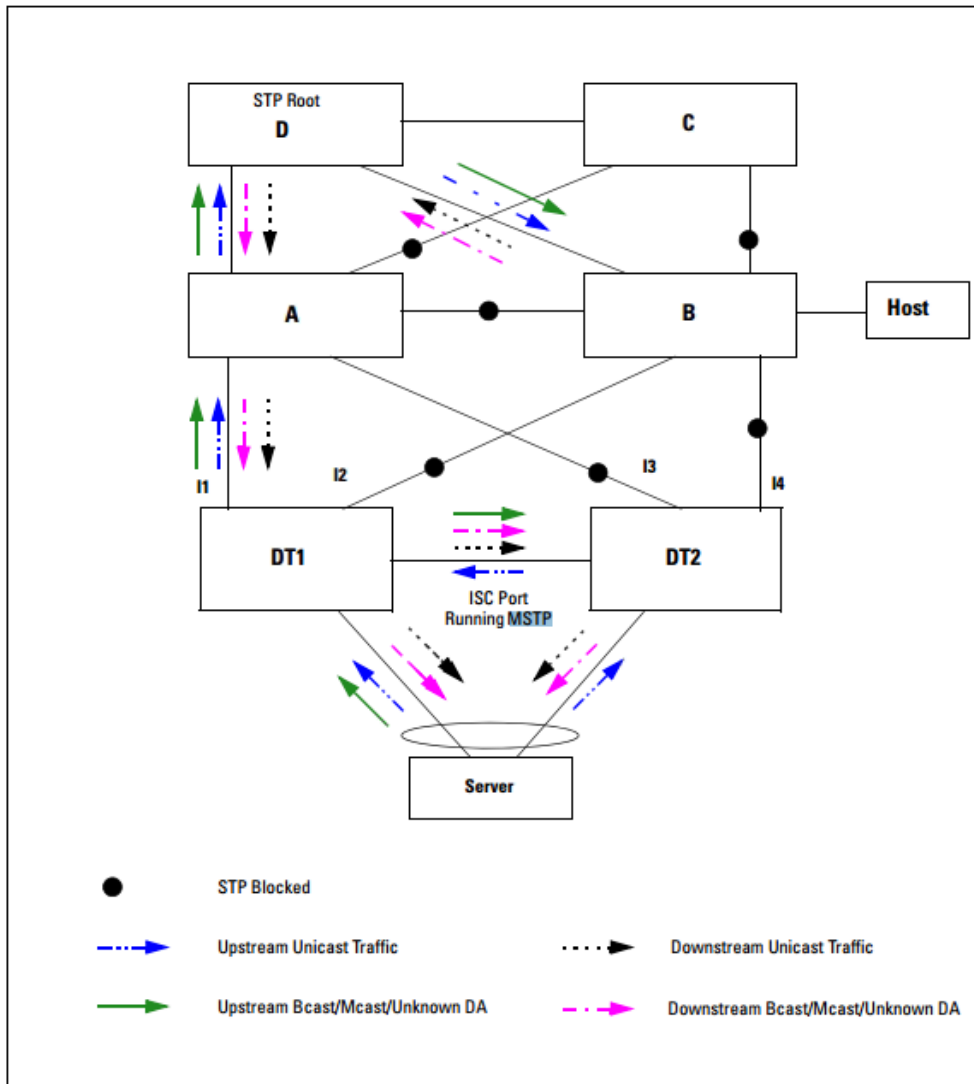


Figure 4-31. Example of Distributed Trunking with STP Forwarding Unicast, Broadcast and Multicast Traffic

C:\Users\1\Google Drive\\_Good\_job\odesk\HP FlexCore\1.png

Reference: HP Switch Software, Management and Configuration Guide

[http://h20628.www2.hp.com/km-ext/kmcsdirect/emr\\_na-c03015912-2.pdf](http://h20628.www2.hp.com/km-ext/kmcsdirect/emr_na-c03015912-2.pdf)

**Question No : 8**

HP Comware Switch 1 connects to switch 2 on GigabitEthernet 1/0/1. Switch 2 implements an inbound rate limit of 600 Mbps. The network administrator wants switch 1 to buffer traffic that exceeds the Switch 2 rate limit of 600 Mbps and send the traffic at 600 Mbps. All traffic