

Exam : F5 Networks F50-521

Title : F5 BIG-IP V9.4 LTM Essentials

Update : DEMO

<http://www.test4actual.com>

Part: A**1:**

Where is the load-balancing mode specified?

A.

within the pool definition

B.

within the node definition

C.

within the virtual server definition

D.

within the pool member definition

Answer: A

2:

Which three methods can be used for initial access to a BIG-IP system? (Choose three.)

A.

CLI access to the serial console port

B.

SSH access to the management port

C.

SSH access to any of the switch ports

D.

HTTP access to the management port

E.

HTTP access to any of the switch ports

F.

HTTPS access to the management port

G.

HTTPS access to any of the switch ports

Answer: A B F

3:

A site needs a virtual server that will use an iRule to parse HTTPS traffic based on HTTP header values. Which two profile types must be associated with such a virtual server? (Choose two.)

A.

TCP

B.

HTTP

C.

HTTPS

D.

ServerSSL

Answer: A B

4:

Which statement accurately describes the difference between two load-balancing modes specified as "member" and "node"?

A.

There is no difference; the two terms are referenced for backward compatibility purposes.

B.

When the load-balancing choice references "node", priority group activation is unavailable.

C.

Load-balancing options referencing "nodes" are available only when the pool members are defined for the "any" port.

D.

When the load-balancing choice references "node", the addresses' parameters are used to make the load-balancing choice rather than the member's parameters.

Answer: D

5:

Which two statements describe differences between the active and standby systems? (Choose two.)

A.

Monitors are performed only by the active system.

B.

Fail-over triggers only cause changes on the active system.

C.

Virtual server addresses are hosted only by the active system.

D.

Configuration changes can only be made on the active system.

E.

Floating self-IP addresses are hosted only by the active system.

Answer: C E**6:**

A site wishes to perform source address translation on packets from some clients but not others. The determination is not based on the client's IP address, but on the virtual server their packets arrive on. What could accomplish this goal?

A.

A SNAT for all addresses could be defined, and then disable the SNAT processing for select VLANs.

B.

Some virtual servers could be associated with SNAT pools and others not associated with SNAT pools.

C.

The decision to perform source address translation is always based on VLAN. Thus, the goal cannot be achieved.

D.

The decision to perform source address translation is always based on a client's address (or network). Thus, this goal cannot be achieved.

Answer: B**7:**

You have a pool of servers that need to be tested. All of the servers but one should be tested every 10 seconds, but one is slower and should only be tested every 20 seconds. How do you proceed?

- A.**
It cannot be done. All monitors test every five seconds.
- B.**
It can be done, but will require assigning monitors to each pool member.
- C.**
It cannot be done. All of the members of a pool must be tested at the same frequency.
- D.**
It can be done by assigning one monitor to the pool and a different monitor to the slower pool member.

Answer: D

8:

Where is connection mirroring configured?

- A.**
It an option within a TCP profile.
- B.**
It is an optional feature of each pool.
- C.**
It is not configured; it is default behavior.
- D.**
It is an optional feature of each virtual server.

Answer: D

9:

When using the setup utility to configure a redundant pair, you are asked to provide a "Failover Peer IP". Which address is this?

- A.**
an address of the other system in its management network
- B.**

an address of the other system in a redundant pair configuration

C.

an address on the current system used to listen for fail-over messages from the partner BIG-IP

D.

an address on the current system used to initiate mirroring and network fail-over heartbeat messages

Answer: B

10:

Assume a virtual server has a ServerSSL profile. What SSL certificates are required on the BIG-IP?

A.

No SSL certificates are required on the BIG-IP.

B.

The BIG-IP's SSL certificates must only exist.

C.

The BIG-IP's SSL certificates must be issued from a certificate authority.

D.

The BIG-IP's SSL certificates must be created within the company hosting the BIG-IPs.

Answer: A

11:

When upgrading a BIG-IP redundant pair, what happens when one system has been updated but the other has not?

A.

Synching should not be performed.

B.

The first system to be updated will assume the Active role.

C.

This is not possible since both systems are updated simultaneously.

D.

The older system will issue SNMP traps indicating a communication error with the partner.

Answer: A

12:

Assume a client's traffic is being processed only by a NAT; no SNAT or virtual server processing takes place. Also assume that the NAT definition specifies a NAT address and an origin address while all other settings are left at their defaults. If a client were to initiate traffic to the NAT address, what changes, if any, would take place when the BIG-IP processes such packets?

A.

The source address would not change, but the destination address would be translated to the origin address.

B.

The destination address would not change, but the source address would be translated to the origin address.

C.

The source address would not change, but the destination address would be translated to the NAT's address.

D.

The destination address would not change, but the source address would be translated to the NAT's address.

Answer: A

13:

Which statement is true concerning SSL termination?

A.

A virtual server that has both ClientSSL and ServerSSL profiles can still support cookie persistence.

B.

Decrypting traffic at the BIG-IP allows the use of iRules for traffic management, but increases the load on the pool member.

C.

When any virtual server uses a ClientSSL profile, all SSL traffic sent to the BIG-IP is decrypted before it is forwarded to servers.

D.

If a virtual server has both a ClientSSL and ServerSSL profile, the pool members have less SSL processing than if the virtual server had only a ClientSSL profile

Answer: A

14:

You have created a custom profile named TEST2. The parent profile of TEST2 is named TEST1. If additional changes are made to TEST1, what is the effect on TEST2?

A.

All changes to TEST1 are propagated to TEST2.

B.

Some of the changes to TEST1 may propagate to TEST2.

C.

Changes to TEST1 cannot affect TEST2 once TEST2 is saved.

D.

When TEST1 is changed, the administrator is prompted and can choose whether to propagate changes to TEST2.

Answer: B

15:

Assume a BIG-IP has no NATs or SNATs configured. Which two scenarios are possible when client traffic arrives on a BIG-IP that is NOT destined to a self-IP? (Choose two.)

A.

If the destination of the traffic does not match a virtual server, the traffic will be discarded.

B.

If the destination of the traffic does not match a virtual server, the traffic will be forwarded based on routing tables.

C.

If the destination of the traffic matches a virtual server, the traffic will be processed per the virtual servers definition.

D.

If the destination of the traffic matches a virtual server, the traffic will be forwarded, but it cannot be load-balanced since no SNAT has been configured.

Answer: A C

16:

If a client's browser does not accept cookies, what occurs when the client connects to a virtual server using cookie persistence?

- A.**
The connection request is not processed.
- B.**
The connection request is sent to an apology server.
- C.**
The connection request is load-balanced to an available pool member.
- D.**
The connection request is refused and the client is sent a "server not available" message.

Answer: C

17:

Assuming other fail-over settings are at their default state, what would occur if the fail-over cable were to be disconnected for five seconds and then reconnected?

- A.**
As long as network communication is not lost, no change will occur.
- B.**
Nothing. Fail-over due to loss of voltage will not occur if the voltage is lost for less than ten seconds.
- C.**
When the cable is disconnected, both systems will become active. When the voltage is restored, unit two will revert to standby mode.
- D.**
When the cable is disconnected, both systems will become active. When the voltage is restored, both systems will maintain active mode.

Answer: C

18:

Assuming there are open connections through an active system's virtual servers and a fail-over occurs, by default, what happens to the connections?

- A.**
All open connections are lost.

B.

All open connections are maintained.

C.

When persistence mirroring is enabled, open connections are maintained even if a fail-over occurs.

D.

Long-lived connections such as Telnet and FTP are maintained, but short-lived connections such as HTTP are lost.

E.

All open connections are lost, but new connections are initiated by the newly active BIG-IP, resulting in minimal client downtime.

Answer: A**19:**

Which VLANs must be enabled for a SNAT to perform as desired (translating only desired packets)?

A.

The SNAT must be enabled for all VLANs.

B.

The SNAT must be enabled for the VLANs where desired packets leave the BIG-IP.

C.

The SNAT must be enabled for the VLANs where desired packets arrive on the BIG-IP.

D.

The SNAT must be enabled for the VLANs where desired packets arrive and leave the BIG-IP.

Answer: C**20:**

When can a single virtual server be associated with multiple profiles?

A.

Never. Each virtual server has a maximum of one profile.

B.

Often. Profiles work on different layers and combining profiles is common.

C.

Rarely. One combination, using both the TCP and HTTP profile does occur, but it is the exception.

D.

Unlimited. Profiles can work together in any combination to ensure that all traffic types are supported in a given virtual server.

Answer: B

www.test4actual.com