



F5 Networks F5CAB2 Practice Questions

Shared by Shelton on 09-10-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

and their status/statistics]

To increase available bandwidth of an existing Trunk, the BIG-IP Administrator is adding additional interfaces. Which command should the BIG-IP Administrator run from within bash shell?

Options:

- A- `tmsh modify /sys trunk trunk_A interfaces add {1.3 1.4}`
- B- `tmsh create /net trunk trunk_A interfaces add {1.3 1.4}`
- C- `tmsh create /sys trunk trunk_A interfaces add {1.3 1.4}`
- D- `tmsh modify /net trunk trunk_A interfaces add {1.3 1.4}`

Answer:

D

Explanation:

Configuring networking objects in BIG-IP requires using the correct TMSH module path and verb.

Module Path: Trunks are networking objects, so they reside under the /net module, not /sys.

Verb: Since the trunk already exists, the modify verb must be used to update its properties rather than create.

Syntax: The correct syntax to append physical interfaces to an existing trunk is `tmsh modify /net trunk [name] interfaces add { [port list] }`.

Data Plane Impact: Adding interfaces to a trunk (Link Aggregation) increases the total aggregate bandwidth available to the system and provides additional hardware redundancy.

Question 2

Question Type: MultipleChoice

and their status/statistics]

To increase the available bandwidth of an existing trunk, the BIG-IP Administrator plans to add additional interfaces. Which command should the BIG-IP Administrator run from within the bash shell? (Choose one answer)

Options:

- A- `tmsh modify /net trunk trunk_A interfaces add {1.3 1.4}`
- B- `tmsh modify /sys trunk trunk_A interfaces add {1.3 1.4}`
- C- `tmsh create /net trunk trunk_A interfaces add {1.3 1.4}`
- D- `tmsh create /sys trunk trunk_A interfaces add {1.3 1.4}`

Answer:

A

Explanation:

In BIG-IP, a trunk is a Layer 2 network object used to aggregate multiple physical interfaces into a single logical link. This aggregation provides increased bandwidth and link resiliency, commonly in conjunction with LACP.

Key concepts that apply here:

Trunks are managed under the `/net trunk tmsh` hierarchy

Physical interfaces are added or removed using the `modify` command

The `create` command is used only when defining a brand-new trunk, not when updating an existing one

Because the trunk already exists and the goal is to add interfaces, the correct operation is:

`tmsh modify /net trunk trunk_A interfaces add {1.3 1.4}`

This command:

Modifies the existing trunk named `trunk_A`

Adds interfaces 1.3 and 1.4 to the trunk

Immediately increases available bandwidth and redundancy

Why the Other Options Are Incorrect

B uses the `/sys` hierarchy, which is not used for trunks

C attempts to create a trunk that already exists

D uses an incorrect hierarchy and an incorrect operation

Question 3

Question Type: MultipleChoice

What type of Virtual Server is configured with no Pool-members, and proxies traffic to the destination IP address specified by the client device?

Options:

- A- Forwarding (IP)
- B- Standard
- C- Performance (Layer 4)
- D- Stateless

Answer:

A

Explanation:

A Forwarding (IP) virtual server is unique because it does not perform load balancing in the traditional sense.

No Pool Members: Unlike a Standard virtual server, which requires a pool to direct traffic, a Forwarding (IP) virtual server typically has no pool assigned.

Destination-Based Routing: The BIG-IP system looks at the destination IP address in the original packet header sent by the client. It then consults the BIG-IP system's local routing table to determine where to send the packet.

Transparency: It acts as a high-performance router/gateway, often used to forward traffic from internal servers to the internet or across different subnets while still allowing the BIG-IP to apply features like SNAT or bandwidth controllers.

Stateful Tracking: While it forwards traffic based on the routing table, it still creates an entry in the connection table to track the flow (unless it is a Stateless virtual server).

Question 4

Question Type: MultipleChoice

A BIG-IP Administrator assigns the default HTTP health monitor to a pool that has three members listening on port 80. When the administrator connects to each pool member using the curl utility, two of the members respond with a status of 404 Not Found, while the third responds with 200 OK. What will the pool show for member availability? (Choose one answer)

Options:

- A- Two members offline and one member online
- B- Two members online and one member offline
- C- All members offline
- D- All members online

Answer:

A

Explanation:

In BIG-IP LTM, health monitors are used to determine the availability of pool members and directly influence traffic flow decisions in the data plane.

Key characteristics of the default HTTP monitor according to BIG-IP Administration Data Plane Concepts:

Sends an HTTP request (typically GET /)

Expects an HTTP response code of 200 OK

Any response other than 200 is treated as a monitor failure

A failed monitor causes the pool member to be marked offline (down)

In this scenario:

Two pool members return 404 Not Found

A 404 response indicates that the requested object was not found

This does not meet the success criteria of the default HTTP monitor

These two members are therefore marked offline

One pool member returns 200 OK

This matches the expected response

The member is marked online

Resulting Pool Member Availability:

2 members: Offline

1 member: Online

Why the Other Options Are Incorrect:

B -- 404 responses are not considered healthy by the default HTTP monitor

C -- At least one member responds with the expected 200 OK

D -- Members returning 404 responses fail the monitor and cannot be marked online

Key Data Plane Concept Reinforced:

BIG-IP health monitors make binary availability decisions based strictly on configured success criteria. For HTTP monitors, response codes matter---404 is a failure, even if the service is technically reachable.

=====

Question 5

Question Type: MultipleChoice

What type of virtual server should be used to load balance UDP traffic without considering previous connections?

Options:

- A- Forwarding
- B- Standard
- C- Reject
- D- Stateless

Answer:

D

Explanation:

When handling high-volume UDP traffic where the BIG-IP does not need to maintain any session history or relationship between packets, a Stateless virtual server is the appropriate choice.

No Connection Tracking: A stateless virtual server does not create or maintain entries in the BIG-IP connection table. This means the system processes each packet as an individual event, without 'considering previous connections' or packets from the same source.

High Performance: Because the system bypasses the overhead of state management, stateless virtual servers provide the highest possible throughput for UDP and ICMP traffic.

Use Cases: This is ideal for services like DNS (stateless queries) or some types of syslog traffic where each packet is independent and doesn't require the persistence or protocol inspection typically provided by a full-proxy.

Why other options are incorrect:

Forwarding: While a Forwarding (IP) virtual server can handle UDP, it still maintains a state entry in the connection table to ensure return traffic is handled correctly.

Standard: This is a full-proxy virtual server. It is inherently stateful and requires a connection table entry for every flow it manages.

Reject: This is a special virtual server type that simply drops incoming traffic and, in the case of TCP, sends a reset (RST) or, for UDP, sends an ICMP unreachable message. It is not a load balancing type.

Question 6

Question Type: MultipleChoice

The BIG-IP Administrator wants to provide quick failover between the F5 LTM devices that are configured as an HA pair with a single Self IP using the MAC Masquerade feature. The administrator configures MAC masquerade for traffic-group-1 using the following command:

```
`tmsh modify /cm traffic-group traffic-group-1 mac 02:12:34:56:00:00`
```

However, the Network Operations team identifies an issue with using the same MAC address across multiple VLANs. As a result, the administrator enables Per-VLAN MAC Masquerade to ensure a unique MAC address per VLAN by running:

```
`tmsh modify /sys db tm.macmasqaddr_per_vlan value true`
```

What would be the resulting MAC address on a tagged VLAN with ID 1501? (Choose one answer)

Options:

A- 02:12:34:56:01:15

B- 02:12:34:56:dd:05

C- 02:12:34:56:05:dd

D- 02:12:34:56:15:01

Answer:

C

Explanation:

In BIG-IP high availability (HA) configurations, MAC Masquerade is used to speed up failover by allowing traffic-group-associated Self IPs to retain the same MAC address when moving between devices. This prevents upstream switches and routers from having to relearn ARP entries during a failover event, resulting in near-instant traffic recovery.

By default, MAC masquerade applies one MAC address per traffic group, regardless of how many VLANs the traffic group spans. This can create problems in some network designs because the same MAC address appearing on multiple VLANs may violate network policies or confuse switching infrastructure.

To address this, BIG-IP provides Per-VLAN MAC Masquerade, enabled by the database variable:

```
`tm.macmasqaddr_per_vlan = true`
```

When this feature is enabled:

BIG-IP derives a unique MAC address per VLAN

The base MAC address configured on the traffic group remains the first four octets

The last two octets are replaced with the VLAN ID expressed in hexadecimal

The VLAN ID is encoded in network byte order (high byte first, low byte second)

VLAN ID Conversion:

VLAN ID: 1501 (decimal)

Convert to hexadecimal:

1501 = 0x05DD

High byte: 05

Low byte: DD

Resulting MAC Address:

Base MAC: `02:12:34:56:00:00`

Per-VLAN substitution last two bytes = `05:DD`

Final MAC address:

`02:12:34:56:05:dd`

Why the Other Options Are Incorrect:

A (01:15) -- Incorrect hexadecimal conversion of 1501

B (dd:05) -- Byte order reversed (little-endian, not used by BIG-IP)

D (15:01) -- Uses decimal values instead of hexadecimal

Key BIG-IP HA Concept Reinforced:

Per-VLAN MAC Masquerade ensures Layer 2 uniqueness per VLAN while preserving the fast failover benefits of traffic groups, making it the recommended best practice in multi-VLAN HA deployments.

P2P
exams

P2P
exams

To Get Premium Files for F5CAB2 Visit

<https://www.p2pexams.com/products/f5cab2>

For More Free Questions Visit

<https://www.p2pexams.com/f5-networks/pdf/f5cab2>

20%
DISCOUNT

P2P
exams