DUMPSARENA

BIG-IP Administration Data Plane Concepts () exam

F5 F5CAB2

Version Demo

Total Demo Questions: 8

Total Premium Questions: 87

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

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)

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:

Two members offline and one member online is correct. The default BIG-IP HTTP monitor sends an HTTP request to the pool member and evaluates the HTTP response as part of its availability decision. With the default monitor behavior, a successful HTTP response is in the successful status-code range; a **200 OK** response satisfies that requirement. Therefore, the pool member returning 200 OK remains available to the pool and is eligible to receive traffic.

The two members returning **404 Not Found** do not satisfy the default HTTP monitor's successful-response criteria. BIG-IP consequently records monitor failures for those member endpoints and marks them down after the configured monitor failure conditions are met. Pool-member availability is evaluated per member, rather than requiring every member to return the same result. As a result, the pool retains one active member while the two endpoints producing 404 responses are unavailable for load balancing.

Administrators can customize an HTTP monitor's send and receive strings when an application intentionally uses a different URI or expected content. In this case, however, the question specifies the default HTTP monitor, so the response from the root request determines the outcome. See the F5 tmsh HTTP monitor reference and the F5 HTTP monitor documentation.

QUESTION NO: 2

A pool member has the following monitor rule assigned:

```
min 2 of { tcp_monitor http_monitor https_monitor }
```

The TCP and HTTP monitors succeed, but the HTTPS monitor fails. What availability state will BIG-IP report for the pool member? (Choose one answer)

A. The pool member remains available because two monitors are successful.

B. The pool member is unavailable because the HTTPS monitor failed.

C. The pool member is unavailable because all assigned monitors must succeed.

D. The pool member status is unknown until all three monitors succeed.

ANSWER: A

Explanation:

The pool member remains available is correct. A monitor rule using `min 2 of` requires at least two monitors in the specified set to succeed. In this case, both the TCP and HTTP monitors are successful, satisfying the required minimum even though the HTTPS monitor fails.

The member would be unavailable if fewer than two monitors succeeded. BIG-IP does not require every monitor in the rule to succeed when the monitor rule explicitly uses a minimum-success threshold.

QUESTION NO: 3

What type of virtual server will have a destination IP address of 0.0.0.0 and listen on a specific VLAN for requests?

A. Forwarding (IP)

B. Wildcard

C. Forwarding (Layer 2)

D. Standard

ANSWER: A

Explanation:

Forwarding (IP) is the correct virtual-server type because it is designed to accept IP traffic that reaches BIG-IP without requiring clients to target a unique, application-specific virtual address. It is normally configured with the wildcard destination address **0.0.0.0** and an any-port setting, allowing BIG-IP to process traffic for destinations not represented by a more-specific virtual server. The virtual server can be restricted through its VLAN and Traffic Management Microkernel (TMM) settings so that it accepts requests arriving on the intended VLAN. BIG-IP then forwards the packets at Layer 3 according to the system routing table, rather than selecting a load-balancing pool as the fundamental purpose of the virtual server. This makes the type useful when BIG-IP is deployed inline and must provide routed forwarding behavior for traffic entering a particular network segment. The all-zero destination is therefore a catch-all match used by the IP forwarding configuration, while the VLAN constraint identifies the interface context from which traffic is accepted. F5 documents Forwarding (IP) as a virtual-server type for forwarding IP traffic and documents VLAN controls as part of virtual-server traffic processing. See the F5 virtual server types documentation and the F5 virtual server traffic-processing documentation.

QUESTION NO: 4

A sync-failover device group contains two traffic groups so that application services can be distributed across the BIG-IP devices.

Which two statements describe traffic-group behavior? (Choose two.)

A. Floating self IP addresses and virtual addresses in the same traffic group move together during a traffic-group failover.

B. Different traffic groups can be active on different devices in the same device group.

C. Non-floating self IP addresses move when their device's traffic group fails over.

D. Traffic-group assignment automatically determines ConfigSync device-group membership.

E. A floating object can be active simultaneously on both devices for the same traffic group.

ANSWER: A B

Explanation:

A traffic group is a failover unit for floating traffic-processing objects. A floating self IP address and a virtual address associated with the same traffic group are owned by the active device for that traffic group and move together when ownership changes.

Different traffic groups can be active on different devices in the same device group, which supports an active-active design. Non-floating self IP addresses are device-specific and do not move with traffic-group ownership. ConfigSync device-group membership is configured independently of traffic-group assignment.

QUESTION NO: 5

Which three iRule events are likely to be seen in iRules designed to select a pool for load balancing? (Choose three.)

A. CLIENT_DATA

B. SERVER_DATA

C. HTTP_REQUEST

D. HTTP_RESPONSE

E. CLIENT_ACCEPTED

F. SERVER_SELECTED

G. SERVER_CONNECTED

ANSWER: A C E

Explanation:

`CLIENT_ACCEPTED`, `CLIENT_DATA`, and `HTTP_REQUEST` are appropriate events for iRules that choose a pool because they occur while BIG-IP is processing client-side traffic and before a server-side connection is established for the selected pool member. In `CLIENT_ACCEPTED`, an iRule can immediately make a routing decision using connection metadata such as the client address, local address, or destination port. This is useful when the required selection criteria are available as soon as the client connection is accepted.

`CLIENT_DATA` supports pool selection for protocols in which the relevant routing information is contained in the client payload. An iRule can collect and inspect the initial data, then use the `pool` command to direct the flow to the appropriate pool. For HTTP traffic, `HTTP_REQUEST` is a common selection point because BIG-IP has parsed the request headers, allowing selection based on values such as the host name, URI, HTTP method, or cookies. F5 documents the `pool` command for specifying the pool used for connection load balancing and documents these client-side and HTTP iRule events in its iRules reference. See F5 iRules pool command and F5 iRules event reference.

QUESTION NO: 6

A pool uses Priority Group Activation with a minimum active-member setting of `2`. Three available pool members have priority group `10`, and two available pool members have priority group `5`.

Which two statements are true? (Choose two.)

A. Priority group 5 members are not selected while at least two priority group 10 members are available.

B. Priority group 5 members become eligible if only one priority group 10 member remains available.

C. All available members are selected regardless of priority group.

D. Priority group 10 members receive health-monitor probes more frequently.

E. Priority group 5 members are used only when no priority group 10 members are available.

ANSWER: A B

Explanation:

Priority Group Activation makes BIG-IP use the highest priority group that has at least the configured minimum number of available members. Because priority group 10 currently has three available members, it satisfies the minimum active-member setting of 2 and is the only eligible priority group for new load-balancing selections.

If one priority-10 member becomes unavailable, only one member remains in that group. BIG-IP then makes the lower-priority members eligible so that the pool has sufficient active members. Priority groups do not control monitor frequency, and the lower-priority group does not wait until every higher-priority member is unavailable when the minimum active-member threshold is not met.

QUESTION NO: 7

A BIG-IP Administrator has a**cluster of devices**.

What should the administrator do**after creating a new Virtual Server on device 1**? (Choose one answer)

A. Synchronize the settings of the**group to device 1**

B. Create a new cluster on device 1

C. Synchronize the settings of**device 1 to the group**

D. Create a new virtual server on device 2

ANSWER: C

Explanation:

After creating the virtual server on device 1, the administrator should select **Synchronize the settings ofdevice 1 to the group**. A virtual server is part of the BIG-IP traffic-management configuration, and in a high-availability deployment that configuration must be synchronized to the other members of the device group. Pushing the configuration from the device where the change was made ensures that the newly created virtual server, along with its associated configuration objects as applicable, is available consistently on the peer device or devices.

BIG-IP Device Service Clustering uses a sync-failover device group to replicate configuration changes among members. Configuration synchronization can be performed manually after making changes, or automatically if automatic synchronization is enabled for the device group. In either case, the authoritative direction for this change is from device 1—the device on which the administrator created the virtual server—to the group. This maintains configuration parity, which is essential so that a failover device can provide the intended application service if it becomes active. F5 documents the configuration synchronization workflow in its BIG-IP Device Service Clustering documentation and describes device-group synchronization behavior in the F5 configuration synchronization overview.

QUESTION NO: 8

A pool member must be removed for maintenance. Which two statements correctly distinguish the Disabled and Forced Offline administrative states? (Choose two.)

A. A Disabled member can continue to receive connections that match an existing persistence record.

B. A Disabled member immediately terminates all existing connections.

C. A Forced Offline member can receive new persistent connections but not new non-persistent connections.

D. A Forced Offline member does not receive new connections, including persistent connections.

E. A Forced Offline member automatically removes all existing connection-table entries.

ANSWER: A D

Explanation:

A member in the **Disabled** state does not receive new non-persistent connections, but existing connections and connections matching an existing persistence record can continue. This state supports a gradual maintenance drain while preserving established client affinity.

A member in the **Forced Offline** state does not receive new connections, including connections that would otherwise be selected because of persistence. Forced Offline does not itself remove connections that already exist in the connection table; those connections must be removed separately when immediate termination is required.