PassQueen
a.k.a dumpsqueen.com

BIG-IP Administration Data Plane Configuration () exam

F5 F5CAB3

Version Demo

Total Demo Questions: 10

Total Premium Questions: 103

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

A pool member supports a resource-intensive application. The BIG-IP Administrator must protect that individual member from both excessive concurrent connections and sudden bursts of new connections.

Which two pool member settings should the BIG-IP Administrator configure? (Pick the 2 correct responses below)

A. Connection Limit

B. Connection Rate Limit

C. Slow Ramp Time

D. Load Balancing Method

E. Availability Requirement

ANSWER: A B

Explanation:

Connection Limit controls the maximum number of concurrent connections that BIG-IP sends to a pool member. **Connection Rate Limit** controls the rate of new connections sent to that member. Together, these settings address both stated capacity constraints.

Slow Ramp Time gradually increases traffic to a newly available member, but it is not a fixed concurrent-connection or new-connection-rate limit. The pool load-balancing method selects among eligible members but does not itself enforce either threshold.

QUESTION NO: 2

A BIG-IP Administrator is creating a trunk to connect to a switch that uses LACP. The BIG-IP system must actively negotiate the link aggregation with the switch.

Which two trunk settings should the BIG-IP Administrator configure? (Pick the 2 correct responses below)

A. Enable LACP on the trunk

B. Set the LACP mode to Active

C. Configure static link aggregation without LACP

D. Set the LACP mode to Passive

E. Create a VLAN group for the trunk interfaces

ANSWER: A B

Explanation:

The trunk must have **LACP enabled** so that it participates in the Link Aggregation Control Protocol. The administrator should also configure **LACP mode Active**, causing BIG-IP to actively send LACP control packets and negotiate with the switch.

Static link aggregation does not use LACP negotiation. Passive mode waits for LACP packets from the peer and does not meet the stated requirement for BIG-IP to actively negotiate. A VLAN group is unrelated to physical link aggregation.

QUESTION NO: 3

The BIG-IP Administrator has configured an HTTP health monitor applied to a Pool of HTTP web servers hosting but all Pool Members show a DOWN status. The web server is returning a response of ' 400 Bad Request ' . What would be the correct monitor Send string?

A. GET /healthmonitor.html HTTP/1.1 \r\nConnection: Close\r\n

B. GET /healthmonitor.html HTTP/1.1 Connection: Close

C. GET /healthmonitor.html HTTP/1.1 \r\nHost: www.f5.com\r\nConnection: Close\r\n\r\n

D. GET /healthmonitor.html HTTP/1.1\r\nConnection: Close\r\n\r\n

ANSWER: C

Explanation:

GET /healthmonitor.html HTTP/1.1 \r\nHost: www.f5.com\r\nConnection: Close\r\n\r\n is the correct monitor Send string because it forms a complete HTTP/1.1 request and supplies the mandatory `Host` request header. A BIG-IP HTTP monitor sends the configured string directly to each pool member; therefore, the request must be valid for the web server and virtual host being monitored. When HTTP/1.1 is used, the server needs the Host value to identify the intended target site, particularly when a server hosts more than one website on the same IP address. The value `www.f5.com` must match the hostname the monitored application expects.

The literal `\r\n` sequences provide the required carriage-return/line-feed delimiters between the request line and headers. The final `\r\n\r\n` terminates the HTTP header section, allowing the server to process the request properly. `Connection: Close` requests that the server close the connection after its response, which is appropriate for a simple health-check transaction. BIG-IP then evaluates the response using the monitor's receive string or standard response handling. HTTP/1.1 specifies that a client must send a Host header in an HTTP/1.1 request; see RFC 9112, section 3.2. F5 documents the HTTP monitor Send field and its request-string behavior in the BIG-IP HTTP monitor API reference.

QUESTION NO: 4

A BIG-IP Administrator configures a node with a standard icmp Health Monitor. The Node shows as DOWN although the Backend Server is configured to answer ICMP requests. Which step should the administrator take next to find the root cause of this issue?

A. Run a qkview

B. Run a tcpdump

C. Run an ssldump

D. Run a curl

ANSWER: B

Explanation:

Run a tcpdump is the appropriate next diagnostic step because an ICMP monitor depends on successful request and reply packets between the BIG-IP system and the monitored node. A packet capture taken on the relevant BIG-IP interface or VLAN shows whether the system is transmitting ICMP echo requests, whether echo replies return, and which source and destination addresses are actually used. This immediately helps isolate the problem domain: monitor traffic may not be leaving through the expected route, a firewall or network device may be filtering ICMP, return traffic may be following an asymmetric path, or the node may be replying to a different source address than expected. Capturing the actual packets is particularly valuable because a successful manual ping from another host does not prove that the BIG-IP monitor's traffic has the same path, source address, routing behavior, or filtering policy. Use a focused capture for the node address and ICMP protocol, then compare transmitted requests with received replies and their timing. F5 documents tcpdump as the primary packet-tracing utility for examining traffic processed by BIG-IP: F5: Overview of packet tracing with tcpdump. Monitor status and monitor behavior are also described in the BIG-IP health monitors documentation.

QUESTION NO: 5

Which two load balancing methods consider all the connections the BIG-IP has between it and each backend application server (Pool Member) when making a load balancing decision for a new connection?

A. Least Connections (node)

B. Ratio (member)

C. Round Robin

D. Weighted Least Connections (node)

ANSWER: A D

Explanation:

Least Connections (node) and **Weighted Least Connections (node)** are correct because their decisions use connection statistics at the node level rather than only at the pool-member level. A node represents the backend server address, so node-level connection counting accounts for active BIG-IP connections to that server across pool members that resolve to the same node. This gives BIG-IP a server-wide view of connection load, which is especially important when one application server provides multiple services on different ports or belongs to multiple pools.

With **Least Connections (node)**, BIG-IP selects the eligible node with the lowest active connection count. With **Weighted Least Connections (node)**, BIG-IP still evaluates the node-wide active connection count, while also applying the configured weighting so that higher-capacity servers can receive a proportionally greater share of connections. In both cases, the defining behavior is that the calculation includes all applicable connections between BIG-IP and the backend node, not merely those associated with one specific member endpoint. F5 documents node-based and member-based load-balancing methods in its BIG-IP Local Traffic Manager Basics documentation and its LB::mode iRules reference.

QUESTION NO: 6

A Virtual Server must accept client traffic only from the VLAN named external. The BIG-IP Administrator must prevent connections arriving on all other VLANs from matching the Virtual Server.

Which two Virtual Server settings should be configured? (Pick the 2 correct responses below)

A. Set VLAN and Tunnel Traffic to Enabled on

B. Add the external VLAN to the enabled VLAN list

C. Set VLAN and Tunnel Traffic to All VLANs and Tunnels

D. Create a VLAN group containing the external VLAN

E. Enable SNAT Automap

ANSWER: A B

Explanation:

The administrator should select **Enabled on** for the Virtual Server VLAN and Tunnel Traffic setting and add the **external** VLAN to the enabled VLAN list. This restricts matching traffic to the specified VLAN.

Using All VLANs and Tunnels would continue to allow traffic from every VLAN. A VLAN group is not required simply to restrict a Virtual Server listener, and source-address translation does not control the ingress VLAN on which a Virtual Server can be matched.

QUESTION NO: 7

Users report that traffic is negatively affected every time a BIG-IP device fails over. The traffic becomes stabilized after a few minutes. What should the BIG-IP Administrator do to reduce the impact of future failovers?

A. Set up Failover Method to HA Order

B. Enable Failover Multicast Configuration

C. Configure MAC Masquerade

D. Configure a global SNAT Listener

ANSWER: C

Explanation:

Configure MAC Masquerade is the appropriate action because it gives a traffic group a shared, locally administered virtual MAC address. Floating self IP addresses and virtual addresses associated with that traffic group are advertised using this masquerade MAC while the traffic group is active. Following failover, the newly active BIG-IP device uses the same configured MAC address rather than its own interface MAC address. This minimizes the Layer 2 relearning required in the connected network and reduces disruption caused by stale forwarding or ARP information during the active-device transition.

MAC masquerading is especially useful where the observed recovery period suggests that switches, routers, firewalls, or other adjacent devices are slow to update forwarding state after floating addresses move. The BIG-IP still performs the required traffic-group ownership transition and sends relevant network announcements, but retaining a consistent source MAC for the floating traffic group helps surrounding infrastructure continue forwarding traffic to the correct active unit more quickly. The masquerade MAC must be unique on the relevant Layer 2 network and should be configured consistently for the traffic group. F5 documents traffic-group MAC masquerading and its failover behavior in the BIG-IP Device Service Clustering documentation and the associated MAC masquerade configuration guidance.

QUESTION NO: 8

Refer to the exhibit.

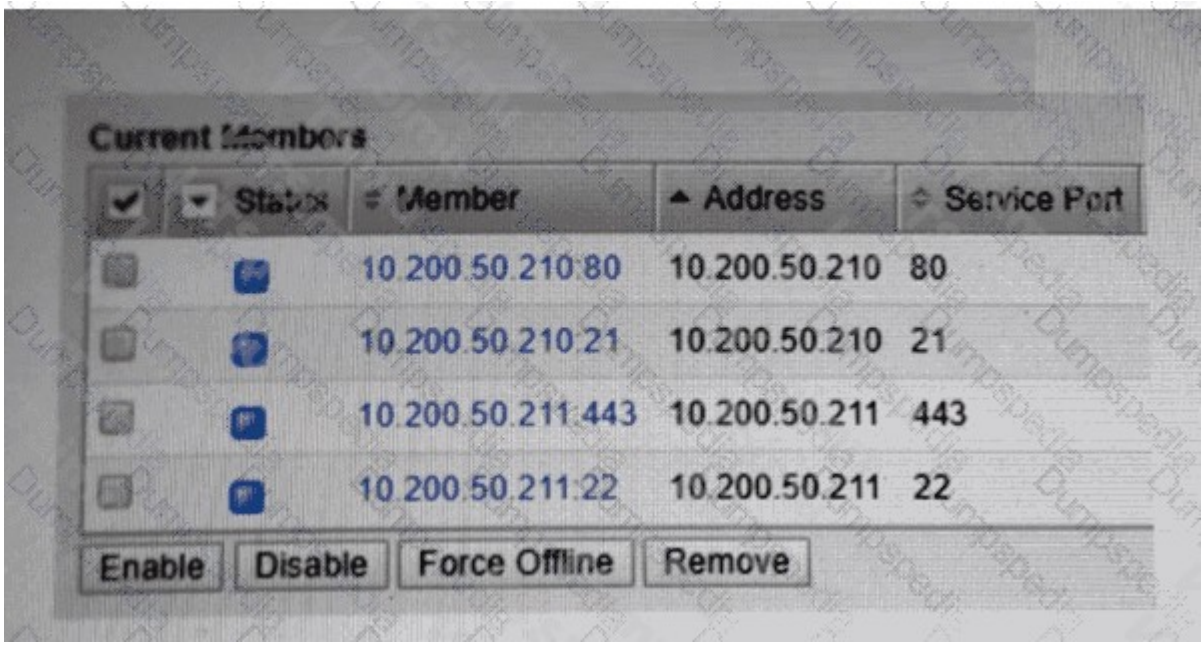

A BIG-IP Administrator needs to configure health monitors for a newly configured server pool named Pool_B.

Which health monitor settings will ensure that all pool members will be accurately marked as available or unavailable? (Choose one answer)

A. HTTPS, HTTP, FTP, and SSH with the Availability Requirement of all health monitors

B. HTTP, HTTPS, FTP, and ICMP with the Availability Requirement of at least one health monitor

C. HTTPS, HTTP, FTP, and SSH with the Availability Requirement of at least one health monitor

D. HTTPS, HTTP, FTP, and SSH with the Availability Requirement of all health monitors

ANSWER: C

Explanation:

HTTPS, HTTP, FTP, and SSH with the Availability Requirement of at least one health monitor is correct because Pool_B contains pool members that provide different application services on their respective service ports. A service-specific monitor is needed for each protocol so that BIG-IP verifies the availability of the application, rather than merely confirming basic network reachability. The HTTPS, HTTP, FTP, and SSH monitors collectively cover the services represented by the pool members.

With an availability requirement of at least one monitor, a member is marked available when the monitor appropriate to its configured service succeeds. This is particularly important for a mixed-service pool: each member does not need to satisfy monitors intended for unrelated application protocols. BIG-IP evaluates the pool monitor rule against each member, and the minimum-success requirement allows the applicable service check to determine that member's availability. This configuration therefore supplies application-aware health validation while accommodating members that listen on distinct ports and use distinct protocols.

F5 documents pool monitor assignment and monitor rules in the tmsh ltm pool reference. F5's monitoring documentation also describes health-monitor behavior and status handling.

QUESTION NO: 9

A Virtual Server must terminate HTTPS connections from clients and establish a separate HTTPS connection to the selected pool member.

Which two SSL profiles should the BIG-IP Administrator apply? (Pick the 2 correct responses below)

A. Client SSL profile

B. Server SSL profile

C. HTTP profile

D. SSL persistence profile

E. OneConnect profile

ANSWER: A B

Explanation:

A **Client SSL profile** is required for BIG-IP to accept and terminate the TLS connection initiated by the client. It contains the client-facing certificate, private key, and TLS settings.

A **Server SSL profile** is required for BIG-IP to initiate TLS to the pool member after the client-side connection has been decrypted. An HTTP profile can provide HTTP-layer processing, but it does not encrypt the server-side connection. A Client SSL profile alone performs SSL offload rather than SSL re-encryption.

QUESTION NO: 10

A BIG-IP Administrator configures a Virtual Server to load balance traffic between 50 webservers for an ecommerce website. Traffic is being load balanced using the Least Connections (node) method. The webserver administrators report that customers are losing the contents from their shopping carts and are unable to complete their orders. What should the BIG-IP Administrator do to resolve the issue?

A. Change Default Persistence Profile setting to cookie

B. Change Default Persistence Profile setting to sip_info

C. Change Load Balancing method to Ratio (node)

D. Change Load Balancing method to Ratio (member)

ANSWER: A

Explanation:

Change Default Persistence Profile setting to cookie. Ecommerce applications commonly require session persistence because a customer's cart and related session state can be maintained by the web server that processed the initial request. Although Least Connections (node) selects an appropriate destination for new connections, it does not inherently ensure that later HTTP requests from the same client return to that same server. Consequently, requests associated with one browser session can reach different pool members, preventing the application from locating the expected server-side cart session.

A BIG-IP cookie persistence profile provides the needed affinity for HTTP traffic. BIG-IP can insert a persistence cookie in the server response, and the browser returns that cookie on subsequent requests. BIG-IP uses the cookie value to associate those requests with the originally selected pool member for the configured persistence timeout. This keeps the customer's session consistently directed to the application instance holding its cart data while still allowing initial client sessions to be distributed across the pool. F5 documents cookie persistence as a persistence method for HTTP applications and describes attaching the persistence profile to the virtual server. See the F5 cookie persistence documentation and the F5 persistence reference.