



F5 Networks F5CAB3 Practice Questions

Shared by Barnes on 09-10-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

A BIG-IP Administrator uses backend servers to host multiple services per server. There are multiple virtual servers and pools defined, referencing the same backend servers. Which load balancing algorithm is most appropriate to have an equal number of connections on each backend server?

Options:

- A- Predictive (node)
- B- Least Connections (node)
- C- Predictive (member)
- D- Least Connections (member)

Answer:

B

Explanation:

This question addresses the critical architectural distinction between 'member-based' and 'node-based' load balancing in environments where servers are multi-homed or host multiple virtualized services. On a BIG-IP, a node represents the underlying IP address of a physical or virtual server, while a pool member represents a specific combination of that IP and a service port (e.g., 10.1.1.10:80).

When multiple pools reference the same IP address across different ports or Virtual Servers, using Least Connections (member) (Option D) only balances connections relative to that specific pool. For example, if Pool_A and Pool_B both use Server_X, and Pool_A has 100 connections while Pool_B has only 5, a member-based algorithm for Pool_B only sees the 5 connections and may continue to send traffic there even if Server_X is CPU-saturated by Pool_A.

The Least Connections (node) algorithm (Option B) solves this by tracking the aggregate total of all active connections directed to that specific node (IP address) across every pool on the system. By selecting the node with the absolute lowest total connection count, the BIG-IP ensures a more equitable distribution of work at the hardware resource level. Predictive methods (Options A and C) use a ranking system based on the trend of connection counts over time, but for the specific requirement of maintaining an equal number of current connections on multi-service servers, the direct Least Connections (node) calculation is the standard and most effective procedural choice.

Question 2

Question Type: MultipleChoice

Refer to the exhibit.

The screenshot displays the F5 BIG-IP Administrator interface for configuring a new VLAN named 'External'. The interface is divided into two main sections: 'General Properties' and 'Resources'.

General Properties:

Name	External
Partition / Path	Common
Description	
Tag	4093

Resources:

Interfaces: 1.2 (untagged)

Buttons: Add, Edit, Delete

Configuration: Advanced

Source Check	☐
MTU	1500
MAC Address	00:50:56:9d:3c:47
Fail-safe	☒
Fail-safe Timeout	90
Action	Reboot
Auto Last Hop	Default
CMP Hash	Default
DAG Tunnel	Outer
IPv6 Prefix Length	128
DAG Adjustment	None

A BIG-IP Administrator configures a new VLAN on an HA pair of devices that does NOT yet have any traffic. This action causes the assigned traffic group to fail over to the standby device. Which VLAN setting should be changed to prevent this issue?

Options:

- A- Auto Last Hop
- B- Source Check
- C- Fail-safe
- D- Customer Tag

Answer:

C

Explanation:

The exhibit shows the advanced configuration of a VLAN where the Fail-safe option is checked. VLAN Fail-safe is a high-availability feature used to monitor network connectivity on a specific VLAN. When enabled, the BIG-IP system monitors the VLAN for network traffic. If the system does not detect any 'useful' traffic on the VLAN within the specified Fail-safe Timeout (which is 90 seconds in the exhibit), it attempts to generate traffic by pinging the default gateway or other devices. If it still detects no traffic, the BIG-IP concludes that the VLAN is unreachable or the network interface has failed, and it triggers a 'Fail-safe Action'---in this case, 'Reboot' or a failover to the peer device in the HA group.

Because the administrator has just created a new VLAN that 'does NOT yet have any traffic,' the Fail-safe mechanism triggers immediately after the 90-second timeout period. Since no devices are yet communicating on this VLAN, the BIG-IP incorrectly assumes there is a hardware or cabling failure and forces a failover to ensure the standby device (which might have better connectivity) takes over. To prevent this unwanted failover, the administrator should uncheck the Fail-safe box for that specific VLAN until the VLAN is fully populated with active nodes and regular traffic. Once the application is live and traffic is flowing, Fail-safe can be re-enabled to provide an additional layer of redundancy. Auto Last Hop (Option A) and Source Check (Option B) are routing and security features that do not trigger HA failover events.

Question 3

Question Type: MultipleChoice

A Standard Virtual Server reports poor network performance for Internet-based clients.

What configuration should be applied?

Options:

- A- Client TCP: f5-tcp-wan / Server TCP: f5-tcp-lan
- B- Client TCP: f5-tcp-lan
- C- Client TCP: f5-tcp-lan / Server TCP: f5-tcp-wan
- D- Client TCP: f5-tcp-optimized

Answer:

A

Explanation:

WAN TCP profiles are optimized for high latency and packet loss typical of Internet clients, while LAN profiles are ideal for backend servers.

Question 4

Question Type: MultipleChoice

All pool members are online and all other settings are default.

What might alter the load balancing behavior?

Options:

- A- Adding a OneConnect profile
- B- Enabling SNAT Automap
- C- Enabling an HTTP fallback host
- D- Adding a persistence profile

Answer:

D

Explanation:

Persistence overrides load balancing decisions by maintaining client-to-server affinity.

Question 5

Question Type: MultipleChoice

Refer to the exhibit.

Visit Us: Brave-Dumps.com

General Properties	
Name	VS-DNS
Partition / Path	Common
Description	
Type	Standard
Source Address	192.168.100.0/23
Destination Address/Mask	192.168.21.50
Service Port	53 Other: ☐
Notify Status to Virtual Address	☒
Link	None
Availability	☐ Unknown (Enabled) - The children pool member(s) eit
Synccookie Status	Off
State	Enabled

Configuration: Advanced

Protocol	TCP
Protocol Profile (Client)	tcp
Protocol Profile (Server)	(Use Client Profile)

DNS queries from two internal DNS servers are being load-balanced to external DNS servers via a virtual server on a BIG-IP device. The DNS queries originate from:

192.168.10.100

192.168.10.200

and target:

192.168.2.150

All DNS queries destined for the external DNS servers fail.

Which property change should the BIG-IP Administrator make in the Virtual Server to resolve this issue? (Choose one answer)

Options:

- A- Protocol profile (Client) to DNS_OPTIMIZED
- B- Type to Performance (HTTP)
- C- Source Address to 192.168.10.0/24
- D- Protocol to UDP

Answer:

D

Explanation:

DNS traffic is primarily transported using UDP port 53. In the exhibit, the Virtual Server is configured with the Protocol set to TCP, which prevents standard DNS queries from being processed correctly. BIG-IP Virtual Servers must be configured with the correct Layer 4 protocol to match the application traffic they are handling.

According to the BIG-IP Administration: Data Plane Configuration documentation:

The Protocol setting on a Virtual Server defines whether traffic is processed as TCP, UDP, or another supported transport protocol.

Standard DNS queries and responses use UDP, while TCP is only required for DNS zone transfers (AXFR) or exceptionally large responses.

When a DNS Virtual Server is incorrectly configured with TCP, UDP-based DNS queries are dropped, causing all requests to fail.

Why the other options are incorrect:

A . Protocol profile (Client) to DNS_OPTIMIZED A DNS profile enhances DNS functionality but does not correct an incorrect transport protocol configuration.

B . Type to Performance (HTTP) Performance (HTTP) Virtual Servers are designed for HTTP traffic and are not suitable for DNS services.

C . Source Address to 192.168.10.0/24 The existing source IPs already fall within the allowed range, so this setting does not address the failure.

Correct Resolution:

Changing the Protocol to UDP aligns the Virtual Server with standard DNS transport requirements, allowing DNS queries to be successfully processed and load-balanced.

Question 6

Question Type: MultipleChoice

A Standard Virtual Server for a web application is configured with Automap for the Source Address Translation option. The original source address of the client must be known by the backend servers. What should the BIG-IP Administrator configure to meet this requirement?

Options:

- A- The Virtual Server type as Performance (HTTP)
- B- An HTTP profile to insert the X-Forward-For header
- C- A SNAT Pool with the client IP
- D- An HTTP Transparent profile

Answer:

B

Explanation:

SNAT Automap is a common configuration that replaces the client's original source IP address with one of the BIG-IP's self IP addresses. This ensures that the backend servers send return traffic back through the BIG-IP, which is necessary for the ADC to process the traffic correctly. However, a side effect of SNAT is that the backend servers only see the BIG-IP's IP in their logs, losing visibility into the true identity of the client.

To resolve this while still using SNAT for routing purposes, the administrator must configure the BIG-IP to 'pass' the client's IP address at the application layer. This is achieved by using an HTTP Profile with the Insert X-Forwarded-For setting enabled. When this profile is applied to the Virtual Server, the BIG-IP intercepts the HTTP request, adds a header (X-Forwarded-For) containing the client's original IP, and then forwards the modified request to the server. The backend web server can then be configured to read this header and log the original client IP instead of the BIG-IP's SNAT address.

Other options are incorrect for this requirement. Performance (HTTP) (Option A) is a virtual server type optimized for speed but often lacks the full Layer 7 header manipulation capabilities of a Standard Virtual Server. SNAT Pool with the client IP (Option C) is technically impossible as SNAT pools use static, pre-defined IPs. There is no such thing as an HTTP Transparent profile (Option D) in standard BIG-IP administration for this purpose. The X-Forwarded-For header insertion within the HTTP profile is the standard procedural method for maintaining client visibility in SNAT-enabled environments.

Question 7

Question Type: MultipleChoice

A BIG-IP Administrator creates a new Virtual Server to load balance SSH traffic. Users are unable to log on to the servers.

What should the BIG-IP Administrator do to resolve the issue?

Options:

- A- Set Protocol to UDP
- B- Set HTTP Profile to None
- C- Set Source Address to 10.1.1.2
- D- Set Destination Address/Mask to 0.0.0.0/0

Answer:

B

Explanation:

SSH is a TCP Layer 4 protocol. Applying an HTTP profile causes BIG-IP to expect HTTP headers, breaking SSH sessions. Removing the HTTP profile allows raw TCP forwarding.

Question 8

Question Type: MultipleChoice

An organization is reporting slow performance accessing their Intranet website. All employees use a single Proxy Server with a public IP.

What should the BIG-IP Administrator do to fix this issue?

Options:

- A- Change Source Address to proxy IP
- B- Change Load Balancing Method to Least Connections
- C- Change Fallback Persistence Profile to source_addr
- D- Change Default Persistence Profile to cookie

Answer:

D

Explanation:

When multiple users share one source IP, source-address persistence fails. Cookie persistence uniquely identifies users at Layer 7 and ensures correct session handling.

Question 9

Question Type: MultipleChoice

The BIG-IP Administrator has to provide encrypted communication between the users and the virtual server they access. Multiple hostnames are configured in DNS with the same IP address. Which profile type and setting in the profile should be used?

Options:

- A- Client SSL, Client Name
- B- Server SSL, Client Name
- C- Server SSL, Server Name
- D- Client SSL, Server Name

Answer:

D

Explanation:

When a single IP address (Virtual Server) must host multiple hostnames (e.g., site1.com, site2.com) over HTTPS, the BIG-IP must be able to present the correct SSL certificate for each site during the initial handshake. This requirement is handled by a technology called Server Name Indication (SNI).

To implement this on the BIG-IP, the administrator must use a Client SSL profile because the encryption is occurring between the user (Client) and the Virtual Server. Within the Client SSL profile settings, there is a field called Server Name. By creating multiple Client SSL profiles---each with a different 'Server Name' (the hostname) and its corresponding certificate---and attaching them all to the same Virtual Server, the BIG-IP can inspect the SNI extension in the

client's 'Client Hello' packet. It then selects the specific profile that matches the requested hostname.

Option A is incorrect because 'Client Name' is not the standard field name used for SNI hostname matching in the profile. Options B and C refer to Server SSL profiles, which manage the connection between the BIG-IP and the backend servers. While Server SSL profiles also have a 'Server Name' field (used to send SNI to a backend server), the requirement here is to provide encrypted communication for 'users' accessing the virtual server, which definitively requires a Client-side SSL configuration.



To Get Premium Files for F5CAB3 Visit

<https://www.p2pexams.com/products/f5cab3>

For More Free Questions Visit

<https://www.p2pexams.com/f5-networks/pdf/f5cab3>

20%
DISCOUNT

P2P
exams