

Dumps F5 F5CAB3 Vce & Valid F5CAB3 Exam Guide



P.S. Free & New F5CAB3 dumps are available on Google Drive shared by Prep4SureReview: <https://drive.google.com/open?id=1kcIQXHZignjf7NioxeU72VcG9zsf9W6B>

Perhaps you have had such an unpleasant experience about F5CAB3 exam questions you brought in the internet was not suitable for you in actual use, to avoid this, our company has prepared F5CAB3 free demo in this website for our customers, with which you can have your first-hand experience before making your final decision. The content of the free demo is part of the content in our real F5CAB3 Study Guide. And you can see how excellent our F5CAB3 training dumps are!

F5 F5CAB3 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">Apply procedural concepts required to modify and manage virtual servers: This domain covers managing virtual servers including applying persistence, encryption, and protocol profiles, identifying iApp objects, reporting iRules, and showing pool configurations.
Topic 2	<ul style="list-style-type: none">Apply procedural concepts required to modify and manage pools: This domain addresses managing server pools including health monitors, load balancing methods, priority groups, and service port configurations.

>> Dumps F5 F5CAB3 Vce <<

TOP Dumps F5CAB3 Vce - F5 BIG-IP Administration Data Plane Configuration - High Pass-Rate Valid F5CAB3 Exam Guide

The F5 F5CAB3 Certification Exam is one of the valuable credentials that are designed to prove an F5 aspirant's technical expertise. With the BIG-IP Administration Data Plane Configuration (F5CAB3) certificate they can be competitive and updated in the highly competitive market. The F5 Certification Questions offers a great opportunity for beginners and experienced professionals to not only validate their skills but also advance their careers.

F5 BIG-IP Administration Data Plane Configuration Sample Questions (Q18-Q23):

NEW QUESTION # 18

A BIG-IP Administrator needs to configure health monitors for a pool containing HTTP, HTTPS, FTP, and SSH services. Which configuration ensures accurate member status?

- A. All monitors with Availability Requirement = at least one
- B. All monitors with Availability Requirement = all
- C. ICMP + TCP with all
- D. HTTP and HTTPS only

Answer: A

Explanation:

Using "at least one" ensures each member is marked up based on its relevant service monitor.

NEW QUESTION # 19

A virtual server is configured to offload SSL from a pool of backend servers. When users connect to the virtual server, they successfully establish an SSL connection but no content is displayed. A packet trace performed on the server shows that the server receives and responds to the request. What should a BIG-IP Administrator do to resolve the problem? (Choose one answer)

- **A. enable SNAT**
- B. enable Server SSL profile
- C. disable Server SSL profile
- D. disable SNAT

Answer: A

Explanation:

This scenario describes a classic case of asymmetric routing in a "one-arm" or non-gateway deployment.

When a BIG-IP system is configured for SSL offloading, the following traffic flow occurs:

* Client-Side: The client establishes a successful SSL/TLS handshake with the Virtual Server. This explains why the user can "successfully establish an SSL connection."

* Server-Side: The BIG-IP decrypts the traffic and forwards it as plain HTTP to the backend server. The packet trace confirms the server receives the HTTP GET request and responds with the content.

* The Routing Failure: By default, the BIG-IP system preserves the client's original source IP address. If the backend server's default gateway is not the BIG-IP system (or if the server is on the same subnet as the client), the server will attempt to send the response directly back to the client's IP address, bypassing the BIG-IP.

* Stateful Drop: Because the BIG-IP is a Full Proxy, it expects the response to return through its own internal state table to be encrypted and sent back to the client. Since the response bypasses the BIG-IP, the BIG-IP connection eventually times out, and the client receives no data despite the server having sent it.

Solution (SNAT): Enabling Secure Network Address Translation (SNAT), specifically SNAT Auto Map, ensures that the BIG-IP replaces the client's source IP with its own internal self-IP before sending the request to the server. This forces the server to send the response back to the BIG-IP, allowing the BIG-IP to complete the transaction and deliver the content to the user.

NEW QUESTION # 20

Application administrators are reporting that nodes different from those configured in the pool are selected.

The use of an iRule is suspected. How can the BIG-IP Administrator check if an iRule is used for this traffic?

(Pick the 2 correct responses below)

- A. Via TMSH with the `list /ltm rule <iRule>` command.
- B. Via the GUI at the iRule tab for the virtual server.
- **C. Via the GUI at the Resources tab for the virtual server.**
- **D. Via TMSH with the `list /ltm virtual <virtual_server>` command.**

Answer: C,D

Explanation:

To determine if an iRule is influencing traffic for a specific Virtual Server, the administrator must verify the association between the Virtual Server object and any applied scripts. In the BIG-IP Configuration Utility (GUI), this association is found under the Resources tab of the specific Virtual Server. While there is an

"iRules" sub-menu under Local Traffic, checking the Virtual Server's Resources tab is the definitive way to see which specific rules are currently active and in what order they are being processed for that particular traffic flow.

From the Command Line Interface (CLI), the `tmsh list /ltm virtual <virtual_server>` command provides a full text-based output of the virtual server's configuration. If iRules are applied, they will appear within a "rules {

... }" block in the command output. This is more effective than Option A, which only lists the contents of the iRule itself but does not show if or where it is applied. Option C is a common misconception; while some versions of the GUI have reorganized menus, the standard location for managing the association of profiles, policies, and iRules to a Virtual Server remains the "Resources" section. By identifying the applied iRule, an administrator can then review the script logic-often containing commands like pool or node-to see if it is overriding the default pool selection based on specific HTTP headers, URI paths, or client IP addresses.

NEW QUESTION # 21

For a given Virtual Server, the BIG-IP must perform SSL Offload and negotiate secure communication over TLSv1.2 only. What should the BIG-IP Administrator do to meet this requirement?

- A. Configure a custom SSL Profile (Client) and select no TLSv1 in the options list
- B. Configure a custom SSL Profile (Server) with a custom TLSv1.2 cipher string
- C. Configure a custom SSL Profile (Client) with a custom TLSv1.2 cipher string
- D. Configure a custom SSL Profile (Server) and select no TLSv1 in the options list

Answer: C

Explanation:

To fulfill the requirement of "SSL Offload" limited to "TLSv1.2 only," the administrator must focus on the client-side of the connection. SSL Offload means the BIG-IP terminates the encrypted connection from the user, processes the traffic (often as plain text internally), and optionally sends it to the backend. The profile responsible for this termination and the initial negotiation with the client's browser is the Client SSL Profile.

A custom Client SSL Profile must be created because the default clientsslprofile typically allows a broad range of protocols for compatibility (including TLS 1.0, 1.1, and 1.2). To restrict communication specifically to TLS 1.2, the administrator modifies the Ciphers string within the profile. Using a string such as DEFAULT:!SSLv3:!TLSv1:!TLSv1.1 or specifically defining TLSv1.2-only suites ensures that the BIG-IP will reject any handshake attempts from older, less secure protocols.

Server SSL Profiles (Options B and C) are used for the encryption between the BIG-IP and the backend nodes, which is not what is requested here. Simply selecting "no TLSv1" in an options list (Option D) is insufficient and often refers to older versions of the software; the modern and standard way to control protocol negotiation on a BIG-IP is through the precise application of Cipher Strings within the Client SSL profile. This ensures compliance with security standards like PCI-DSS while providing the offloading benefits to the backend infrastructure.

NEW QUESTION # 22

A BIG-IP Administrator creates a new Virtual Server. The end user is unable to access the page. During troubleshooting, the administrator learns that the connection between the BIG-IP system and server is NOT set up correctly. What should the administrator do to solve this issue? (Choose one answer)

- A. Set Address Translation to SNAT and have a self-IP configured in the same subnet as the servers
- B. Set Address Translation to Auto Map, configure a SNAT pool, and have pool members in the same subnet as the servers
- C. Set Address Translation to SNAT and configure a specific translation address
- D. Disable Address Translation

Answer: A

Explanation:

The issue described is a classic symptom of asymmetric routing, which frequently occurs when the BIG-IP system and the back-end servers reside on the same subnet (often referred to as a "one-arm" deployment).

The Routing Problem: By default, the BIG-IP system preserves the original client source IP address when forwarding traffic to a pool member. If the server is in the same subnet as the client or if the server's default gateway is not the BIG-IP, the server will attempt to send its response directly back to the client's IP address, bypassing the BIG-IP.

Stateful Failure: Since the BIG-IP is a Full Proxy, it maintains a state table. Because the response packet never returns through the BIG-IP, the system cannot complete the three-way handshake or manage the application session, resulting in a connection failure for the user.

The Solution (SNAT): Enabling Source Network Address Translation (SNAT) solves this by changing the source IP address of the request to an IP address owned by the BIG-IP (typically a self-IP).

Requirement for Subnet Alignment: To ensure the server sends the response back to the BIG-IP, the translation address must be reachable. By using a self-IP configured in the same subnet as the servers, the BIG-IP ensures that the server sees the request coming from a local "neighbor." The server will then naturally send the response back to that self-IP, allowing the BIG-IP to translate the packet back and forward it to the client.

Why other options are incorrect:

A: Disabling address translation would ensure the server-side traffic uses the client IP, making asymmetric routing inevitable in this scenario.

B: This is technically contradictory; "Auto Map" specifically uses existing self-IPs and does not require or use a "SNAT pool" configuration.

C: While using a specific translation address can work, it does not inherently guarantee the Layer 2/Layer 3 reachability mentioned in the scenario as effectively as ensuring the self-IP is correctly placed in the server's subnet.

