F5CAB3 Latest Braindumps Sheet - F5CAB3 Latest Test Vce

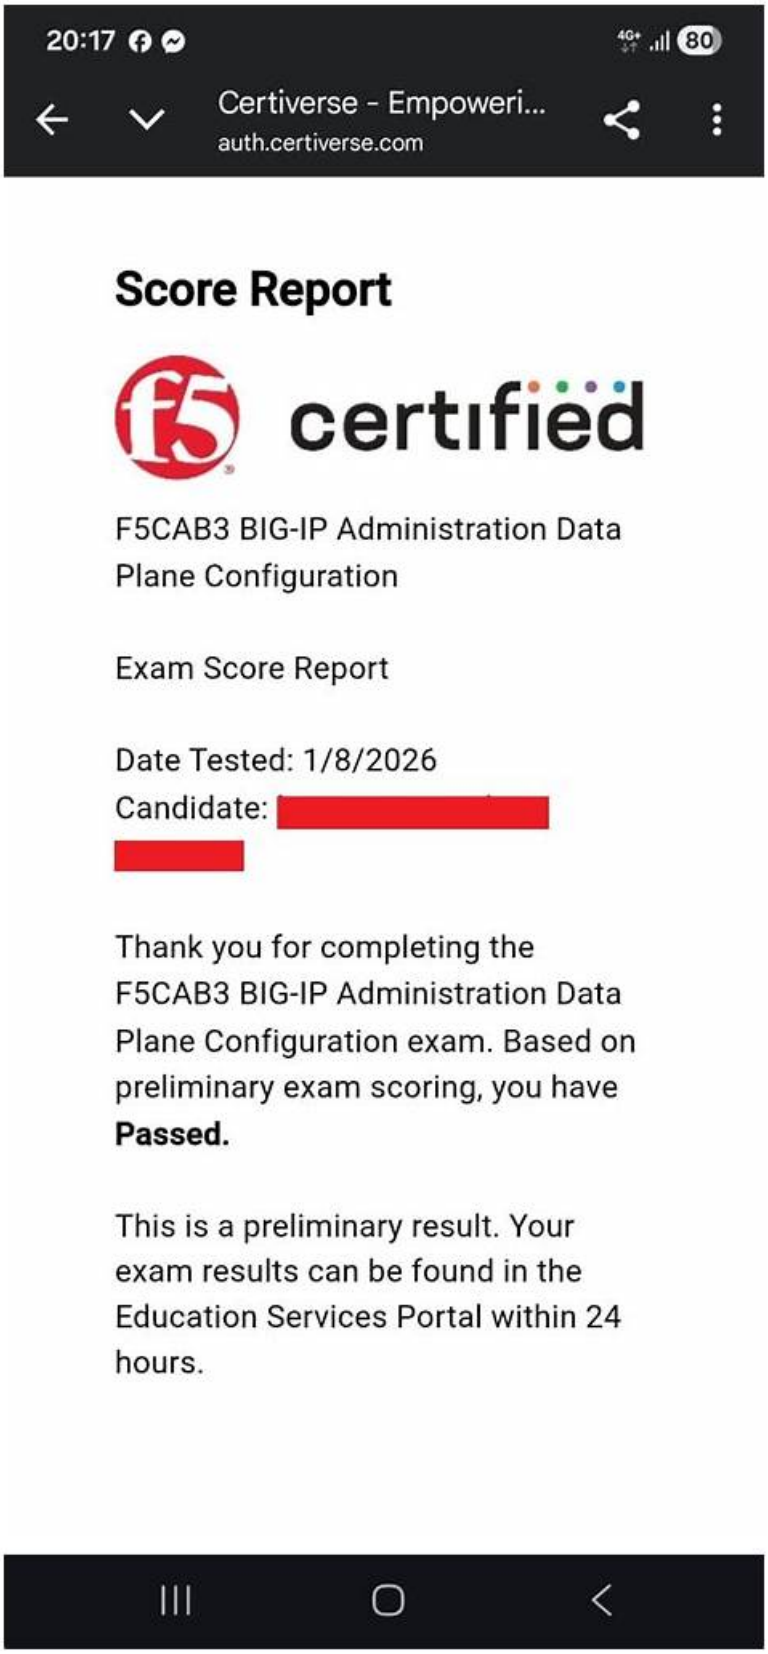

In traditional views, the F5CAB3 practice materials need you to spare a large amount of time on them to accumulate the useful

knowledge may appearing in the real F5CAB3 exam. However, our F5CAB3 learning questions are not doing that way. According to data from former exam candidates, the passing rate of our F5CAB3 learning material has up to 98 to 100 percent. There are adequate content to help you pass the exam with least time and money.

F5 F5CAB3 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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.
Topic 2	• 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.

F5CAB3 Latest Test Vce | F5CAB3 Technical Training

To ensure a more comfortable experience for users of F5CAB3 test material, we offer a thoughtful package. Not only do we offer free demo services before purchase, we also provide three learning modes of F5CAB3 learning guide for users. With easy payment and thoughtful, intimate after-sales service, believe that our F5CAB3 Exam Guide Materials will not disappoint users. Last but not least, our worldwide service after-sale staffs will provide the most considerable and comfortable suggestion on F5CAB3 study prep for you in twenty -four hours a day, as well as seven days a week incessantly.

F5 BIG-IP Administration Data Plane Configuration Sample Questions (Q59-Q64):

NEW QUESTION # 59
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. Round Robin
- C. Ratio (member)
- D. Weighted Least Connections (node)

Answer: A,D

Explanation:
The two load balancing methods that consider all connections between the BIG-IP and each backend node
- not just connections to a specific pool member - are Least Connections (node) and Weighted Least Connections (node) .
The critical distinction here lies in the node-level scope of evaluation. A node represents the backend server ' s IP address, regardless of how many services or ports it may be serving. Therefore:
* Least Connections (node) directs new connections to the node with the fewest total active connections across all services on that server, providing a holistic connection-count perspective.
* Weighted Least Connections (node) operates identically but factors in an administrator-defined ratio weight, allowing servers with greater capacity to proportionally absorb more connections while still evaluating total node-level connection counts.
By contrast:
* Ratio (member) distributes traffic based on a static weight ratio and does not dynamically evaluate current connection counts.
* Round Robin distributes traffic sequentially in rotation, completely ignoring current connection states on any node or member.
The node-based methods are particularly valuable in environments where a single backend server hosts multiple pool members across different ports, ensuring the server ' s overall load - not just per-service load
- governs balancing decisions.
Reference: BIG-IP Administration - Data Plane Configuration, Module: Load Balancing Methods - Member vs. Node Scope.

NEW QUESTION # 60
Refer to the exhibit.

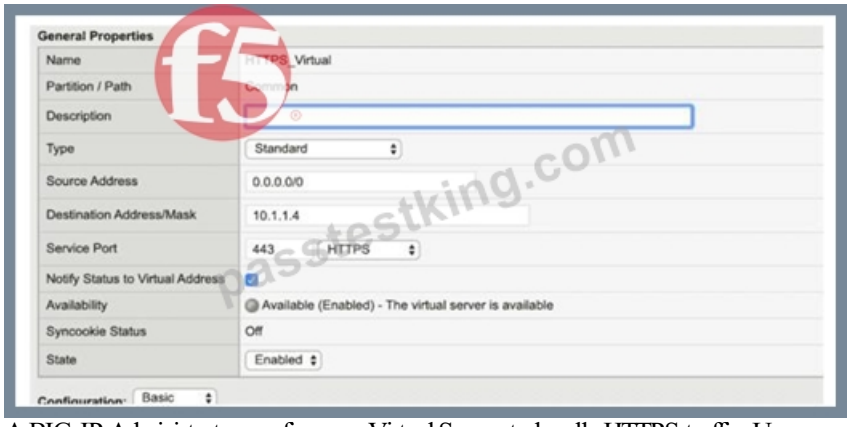

A BIG-IP Administrator configures a Virtual Server to handle HTTPS traffic. Users report that the application is NOT working.
Which additional configuration is required to resolve this issue?

- A. Configure SSL Profile (Server)
- B. Configure Service Port to HTTP
- C. Configure Protocol Profile (Server)
- D. Configure SSL Profile (Client)

Answer: D

Explanation:
According to the provided exhibit, the " SSL Profile (Client) " section in the Virtual Server configuration is empty. For a BIG-IP system to process HTTPS traffic, it must act as an SSL/TLS endpoint. This process, known as SSL Termination or SSL Offload, requires the assignment of a Client SSL Profile to the Virtual Server. Without this profile, the BIG-IP does not have the necessary certificate and private key information to perform the SSL handshake with the client ' s browser. Consequently, when a user attempts to connect via HTTPS, the TCP connection may establish, but the SSL handshake will fail because the BIG-IP will not know how to decrypt the incoming encrypted packets.
A Client SSL profile defines the ciphers, certificates, and keys that the BIG-IP uses to communicate securely with the client. In a standard HTTPS deployment, the BIG-IP decrypts the traffic and can then send it to the backend pool members either as plain text (header insertion/manipulation) or re-encrypt it using a Server SSL profile. While a Server SSL profile (Option C) is needed if the backend servers themselves require HTTPS, the initial failure for a user reaching a Virtual Server is almost always the lack of a Client SSL profile to terminate the user ' s connection. Changing the Service Port to HTTP (Option D) would be incorrect because the goal is to handle HTTPS traffic (typically port 443). Assigning the " clientssl " or a custom client-side profile from the " Available " list to the " Selected " list in the GUI is the mandatory step to make the Virtual Server operational for secure web traffic.

NEW QUESTION # 61
A node is a member of multiple pools and hosts different applications. If one application becomes unavailable, only that pool member should be marked down.
What should the BIG-IP Administrator deploy?

- A. UDP monitor
- B. HTTP monitor with custom send/receive
- C. ICMP + TCP monitor
- D. TCP monitor

Answer: B

Explanation:
Application-level monitors ensure that only the affected service is marked down, not the entire node.

NEW QUESTION # 62
How will the BIG-IP system distribute the traffic based on the configuration below?
pool my_pool {
lb_mode fastest

min_active_members 2
member 10.12.10.7:80 priority 3
member 10.12.10.8:80 priority 3
member 10.12.10.9:80 priority 3
member 10.12.10.4:80 priority 2
member 10.12.10.5:80 priority 2
member 10.12.10.6:80 priority 2
member 10.12.10.1:80 priority 1
member 10.12.10.2:80 priority 1
member 10.12.10.3:80 priority 1
}
(Pick the 2 correct responses below)

- A. Connections are first distributed to all pool members with priority 3 when all the pool members with priority 3 are available
- B. If both the priority 3 group and the priority 2 group have fewer than two members available, traffic is directed to the priority 1 group
- C. Connections are distributed to all pool members with priority 2 if one pool member with priority 3 is down
- D. If both the priority 1 group and the priority 2 group have fewer than two members available, traffic is directed to the priority 3 group

Answer: A,B

Explanation:
The configuration provided utilizes Priority Group Activation in conjunction with the min_active_members setting. Priority groups allow an administrator to define primary servers and "backup" servers within the same pool. The BIG-IP prioritizes traffic based on the assigned priority number, with the highest number receiving traffic first.
In this specific configuration, the priority 3 group is the primary group. Therefore, connections are first distributed to all pool members with priority 3 as long as they are available. The system will continue to use only the priority 3 group unless the number of available members in that group falls below the min_active_members value, which is set to 2.
If the priority 3 group has fewer than two active members, the BIG-IP "activates" the next available priority group (priority 2) and distributes traffic among the remaining members of priority 3 and all members of priority 2. This cascading logic continues down the list. Consequently, if both the priority 3 group and the priority 2 group have fewer than two members available, traffic is directed to the priority 1 group. This ensures that even in a multi-server failure scenario, the system has a last-resort group of servers to handle the traffic.
Option D is incorrect because if only one member of priority 3 goes down, there are still two members active (10.12.10.8 and 10.12.10.9). Since 2 is not less than the min_active_members threshold of 2, the priority 2 group will not yet be activated. Option B is incorrect because traffic flows from high priority to low priority, not the other way around.

NEW QUESTION # 63
What would be the best persistence method for F5 to load balance traffic from clients via a single source IP (NAT) to multiple pool members with even distribution for an HTTPS web application?

- A. Destination address affinity persistence
- B. Source address affinity persistence
- C. Cookie persistence
- D. SSL persistence

Answer: C

Explanation:
When clients connect through a single source IP (NAT) - such as a corporate proxy or carrier-grade NAT
- Source Address Affinity persistence becomes entirely ineffective because all clients share an identical source IP address. This would force every client session to the same pool member, completely eliminating even distribution and defeating the purpose of load balancing.
Cookie Persistence is the optimal solution in this scenario because it operates at Layer 7 , inserting a unique cookie into each client ' s HTTP/HTTPS response. Each individual browser session carries its own distinct cookie value, allowing the BIG-IP to identify and persist individual clients to specific pool members - regardless of whether they share a common source IP address. This guarantees both session persistence and even load distribution across pool members.
The remaining options are unsuitable because:
* Destination Address Affinity persists based on destination IP, irrelevant for client-to-server session stickiness.

* SSL Persistence uses the SSL Session ID for persistence, but SSL session IDs are frequently renegotiated and are unreliable for long-term persistence across modern TLS implementations.
* Source Address Affinity fails entirely under NAT conditions as described.
Cookie persistence is the industry-standard method for HTTPS application load balancing requiring per-client session affinity.
Reference: BIG-IP Administration - Data Plane Configuration, Module: Persistence Methods - Cookie, Source Address, and SSL Persistence.

NEW QUESTION # 64
.....

The simulation of the actual F5 F5CAB3 test helps you feel the real F5CAB3 exam scenario, so you don't face anxiety while giving the final examination. You can even access your last test results, which help to realize your mistakes and try to avoid them while taking the F5 F5CAB3 Certification test.