Free PDF Quiz High Pass-Rate F5CAB2 - Valid BIG-IP Administration Data Plane Concepts (F5CAB2) Study Guide

P.S. Free 2026 F5 F5CAB2 dumps are available on Google Drive shared by ITExamSimulator: https://drive.google.com/open?id=1rjjKSyUA-p3sir-iS8_ipSpWASmkhqRh

In order to help you more ITExamSimulator the F5 F5CAB2 exam eliminate tension of the candidates on the Internet. F5CAB2 study materials including the official F5 F5CAB2 certification training courses, F5 F5CAB2 self-paced training guide, F5CAB2 exam ITExamSimulator and practice, F5CAB2 Online Exam F5CAB2 study guide. F5CAB2 simulation training package designed by ITExamSimulator can help you effortlessly pass the exam. Do not spend too much time and money, as long as you have ITExamSimulator learning materials you will easily pass the exam.

When dealing with any kind of exams, the most important thing is to find a scientific way to review effectively. our F5CAB2 exam materials are elemental materials you cannot miss. For our passing rate of F5CAB2 practice quiz has reached up to 98 to 100 percent up to now. Besides, free updates of F5CAB2 learning guide will be sent to your mailbox freely for one year after payment,and you will have a great experience during usage of our F5CAB2 study prep.

>> Valid F5CAB2 Study Guide <<

Valid Dumps F5CAB2 Questions | Latest F5CAB2 Dumps

Mercenary men lust for wealth, our company offer high quality F5CAB2 practice engine rather than focusing on mercenary motives. They are high quality and high effective F5CAB2 training materials and our efficiency is expressed clearly in many aspects for your reference. The first one is downloading efficiency. The second is expressed in content, which are the proficiency and efficiency of F5CAB2 Study Guide. You will love our F5CAB2 exam questions as long as you have a try!

F5 F5CAB2 Exam Syllabus Topics:

Topic	Details
Topic 1	• Explain the relationship between interfaces, trunks, VLANs, self-IPs, routes and
Topic 2	• their status • statistics: This domain covers BIG-IP networking components including interfaces, trunks, VLANs, self-IPs, and routes, their dependencies and status, plus predicting traffic paths and egress IPs.

Topic 3	• Identify the different virtual server types: This domain covers BIG-IP virtual server types: Standard, Forwarding, Stateless, Reject, Performance Layer 4, and Performance HTTP.

F5 BIG-IP Administration Data Plane Concepts (F5CAB2) Sample Questions (Q18-Q23):

NEW QUESTION # 18
Refer to the exhibit above.
A BIG-IP pool is configured with Priority Group Activation = Less than 2 available members. The pool members have different priority groups and availability states. Which pool members are receiving traffic?
(Choose one answer)

- A. serv1
- B. serv1, serv3
- C. serv1, serv3, serv4
- D. serv1, serv2, serv3, serv4

Answer: C

Explanation:
This question tests understanding of Priority Group Activation (PGA) and how BIG-IP determines which pool members are eligible to receive traffic.
Key BIG-IP Priority Group Concepts:
* Higher priority group numbers = higher priority
* BIG-IP will only send traffic to the highest priority group that meets the Priority Group Activation condition
* Lower priority groups are activated only when the condition is met
* Only available (green) members count toward the activation threshold
Configuration from the Exhibit:
* Priority Group Activation: Less than 2 available members
* Pool Members and Status:
Pool Member
Priority Group
Status
serv1
2
Active (available)
serv2
2
Inactive (down)
serv3
1
Active (available)
serv4
1
Active (available)
Step-by-Step Traffic Decision:
* BIG-IP first evaluates the highest priority group (Priority Group 2)
* Priority Group 2 has:
* serv1 # available
* serv2 # unavailable
* Total available members = 1
* Activation rule is Less than 2 available members
* Condition is true (1 < 2)
* BIG-IP activates the next lower priority group (Priority Group 1)
* Traffic is now sent to:
* serv1 (Priority Group 2)
* serv3 and serv4 (Priority Group 1)
Final Result:

Traffic is distributed to serv1, serv3, and serv4
Why the Other Options Are Incorrect:
* A - Ignores activation of the lower priority group
* B - serv4 is also active and eligible
* C - serv2 is down and cannot receive traffic
Key Data Plane Concept Reinforced:
Priority Group Activation controls when lower-priority pool members are allowed to receive traffic, based strictly on the number of available members in the higher-priority group. In this case, the failure of one high-priority member caused BIG-IP to expand traffic distribution to lower-priority members to maintain availability.

NEW QUESTION # 19
A BIG-IP Administrator wants to add a new Self IP to the BIG-IP device. Which item should be assigned to the new Self IP being configured?

- A. Trunk
- B. VLAN
- C. Interface
- D. Route

Answer: B

Explanation:
ASelf IPis an IP address on the BIG-IP system that you associate with a specificVLAN.
* VLAN Association:A Self IP cannot exist independently; it must be bound to a VLAN to define which network segment the BIG-IP can communicate with.
* Layer 2 to Layer 3 Mapping:While a VLAN is associated with physical interfaces or trunks (Layer 2), the Self IP provides the Layer 3 identity for the BIG-IP on that VLAN.
* Traffic Processing:Self IPs are used by the BIG-IP for health checking backend servers, acting as a default gateway for servers, and for HA heartbeat communication.

NEW QUESTION # 20
A BIG-IP Administrator is informed that traffic on interface 1.1 is expected to increase beyond the maximum bandwidth capacity of the link. There is a single VLAN on the interface.
What should the BIG-IP Administrator do to increase the total available bandwidth? (Choose one answer)

- A. Set the media speed of interface 1.1 manually
- B. Create a trunk object with two interfaces
- C. Increase the MTU on the VLAN using interface 1.1
- D. Assign two interfaces to the VLAN

Answer: B

Explanation:
Comprehensive and Detailed Explanation (BIG-IP Administration - Data Plane Concepts):
On BIG-IP systems, physical interface bandwidth is fixed by the link speed (for example, 1GbE or 10GbE). When traffic demand exceeds the capacity of a single interface, BIG-IP provides link aggregation through trunks.
Key concepts involved:
Interfaces
A single physical interface (such as 1.1) is limited to its negotiated link speed. You cannot exceed this capacity through software tuning alone.
Trunks (Link Aggregation)
A trunk combines multiple physical interfaces into a single logical interface.
BIG-IP supports LACP and static trunks.
Traffic is distributed across member interfaces, increasing aggregate bandwidth and providing redundancy.
VLANs are then assigned to the trunk, not directly to individual interfaces.
Why option B is correct:
Creating a trunk with two interfaces allows BIG-IP to use both physical links simultaneously.
This increases total available bandwidth (for example, two 10Gb interfaces → up to 20Gb aggregate capacity).
This is the documented and supported method for scaling bandwidth on BIG-IP.

Why the other options are incorrect:
A . Increase the MTU
MTU changes affect packet size and efficiency, not total bandwidth capacity.
C . Assign two interfaces to the VLAN
BIG-IP does not support assigning a VLAN to multiple interfaces directly. VLANs must be associated with one interface or one trunk.
D . Set the media speed manually
Media speed can only be set up to the physical capability of the interface and connected switch port. It cannot exceed the hardware limit.
Conclusion:
To increase total available bandwidth on BIG-IP when a single interface is insufficient, the administrator must create a trunk object with multiple interfaces and move the VLAN onto the trunk. This aligns directly with BIG-IP data plane design and best practices.

NEW QUESTION # 21
An organization needs to deploy an HTTP application on a BIG-IP system. The requirements specify hardware acceleration to enhance performance, while HTTP optimization features are not required.
What type of virtual server and associated protocol profile should be used to meet these requirements?
(Choose one answer)

- A. Type: Performance (HTTP) Protocol Profile: fasthttp
- B. Type: Stateless Protocol Profile: fastL4
- C. Type: Performance (Layer 4) Protocol Profile: fastL4
- D. Type: Standard Protocol Profile: tcp-wan-optimized

Answer: C

Explanation:
To select the correct virtual server type, an administrator must balance the need for L7 intelligence versus raw throughput and hardware offloading:
* Performance (Layer 4) Virtual Server: This type is designed for maximum speed. It uses the fastL4 profile, which allows the BIG-IP system to leverage the ePVA (Embedded Packet Velocity Accelerator) hardware chip. When a Performance (L4) virtual server is used, the system processes packets at the network layer (L4) without looking into the application payload (L7). This fulfills the requirement for hardware acceleration and avoids the overhead of HTTP optimization features, which are not needed in this scenario.
* Performance (HTTP) Virtual Server: While fast, this type uses the fasthttp profile to provide some L7 awareness and optimization (like header insertion or small-scale multiplexing). Since the requirement specifically states HTTP optimization is not required, the L4 variant is more efficient.
* Standard Virtual Server: This is a full-proxy type. While it offers the most features (SSL offload, iRules, Compression), it processes traffic primarily in the TMOS software layer (or via high-level hardware assistance), which is "slower" than the pure hardware switching path of the Performance (L4) type.
* Stateless Virtual Server: This is typically used for specific UDP/ICMP traffic where the system does not need to maintain a connection table. It is not appropriate for standard HTTP (TCP) applications requiring persistent sessions or stateful load balancing.
By choosing Performance (Layer 4) with the fastL4 profile, the organization ensures that the traffic is handled by the hardware acceleration chips, providing the lowest latency and highest throughput possible for their HTTP application.

NEW QUESTION # 22
In the GUI, where should the BIG-IP Administrator configure an existing VLAN named external to a specific interface as untagged?

- A. Network - VLANs -> VLAN List -> external -> interfaces
- B. Network -> VLANs -> VLAN List -> internal -> interfaces
- C. Network -> VLANs -> VLAN List -> create -> external

Answer: A

Explanation:
In the BIG-IP system, VLANs are the logical entities that group physical interfaces or trunks together. To modify how a VLAN interacts with an interface, the administrator must navigate to the specific VLAN configuration object.
* VLAN List:This section displays all existing VLANs configured on the system.
* Interface Association:Within the properties of a specific VLAN (in this case, "external"), there is an Interfacessection. This is

where physical ports or trunks are assigned to the VLAN.
* Tagging Status:For each associated interface, the administrator can choose betweenTagged(802.1Q) or Untagged.
* Untagged:The interface will treat incoming traffic without a VLAN header as part of this VLAN, and outgoing traffic will not have a VLAN tag added. An interface can only be "Untagged" for one VLAN.
* Tagged:Allows an interface to carry traffic for multiple VLANs simultaneously by using 802.1Q headers.

NEW QUESTION # 23
......

Our F5CAB2 exam braindump is revised and updated according to the change of the syllabus and the latest development situation in the theory and the practice. The F5CAB2 exam torrent is compiled elaborately by the experienced professionals and of high quality. The contents of F5CAB2 guide questions are easy to master and simplify the important information. It conveys more important information with less answers and questions, thus the learning is easy and efficient. The language is easy to be understood makes any learners have no obstacles to study and pass the F5CAB2 Exam.