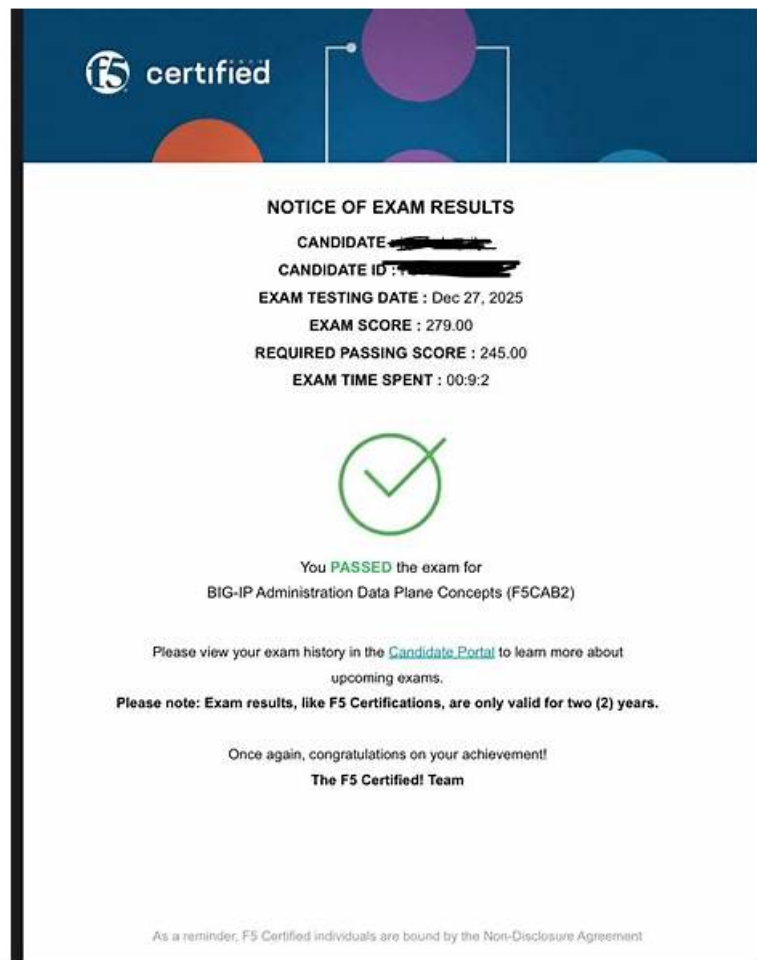


Test F5CAB2 Dumps Pdf - F5CAB2 Latest Dump



BONUS!!! Download part of Fast2test F5CAB2 dumps for free: https://drive.google.com/open?id=1qLrdnFMkohWuajhkYp_VnmG5JH-q465L

Hence, memorizing them will help you get prepared for the F5 F5CAB2 examination in a short time. The product of Fast2test comes in PDF, desktop practice exam software, and BIG-IP Administration Data Plane Concepts (F5CAB2) (F5CAB2) web-based practice test. To give you a complete understanding of these formats, we have discussed their features below.

The three versions of our F5CAB2 training materials each have its own advantage. On the one hand, the software version can simulate the real F5CAB2 examination for all of the users in windows operation system. By actually simulating the real test environment. On the other hand, if you choose to use the software version, you can download our F5CAB2 ExamPrep only for Windows system. We strongly believe that the software version of our F5CAB2 study materials will be of great importance for you to prepare for the exam and all of the employees in our company wish you early success.

>> Test F5CAB2 Dumps Pdf <<

100% Pass Quiz 2026 F5 F5CAB2: BIG-IP Administration Data Plane Concepts (F5CAB2) High Hit-Rate Test Dumps Pdf

Studying with us will help you build the future you actually want to see. By giving you both the skills and exposure of your area of work, our F5CAB2 study guides, F5CAB2 dump and practice questions and answers will help you pass F5CAB2 Certification without any problem. Our very special F5CAB2 products which include F5CAB2 practice test questions and answers encourage you to think higher and build a flourishing career in the every growing industry.

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

(Q10-Q15):

NEW QUESTION # 10

The network architecture for a BIG-IP consists of an external VLAN and an internal VLAN with two interfaces connected to the upstream switch. The design requires fault tolerance in the case that one of the interfaces is down. Which deployment architecture meets these requirements? (Choose one answer)

- A. Two network trunks each with one VLAN and LACP enabled, and both VLANs configured as tagged
- B. Two network trunks each with one VLAN and LACP disabled, and one VLAN configured as tagged and one VLAN configured as untagged
- C. One network trunk with both VLANs and LACP enabled, and both VLANs configured as untagged
- **D. One network trunk with both VLANs and LACP enabled, and both VLANs configured as tagged**

Answer: D

Explanation:

To meet the requirement of fault tolerance when one interface goes down, BIG-IP must use link aggregation so that loss of a single physical link does not isolate the VLAN(s).

How the objects relate (data plane view)

* Interfaces = physical links.

* Trunk (LACP) = bundles multiple interfaces into one logical link that provides redundancy (and possibly bandwidth aggregation).

* VLANs are assigned to interfaces or trunks. If you need multiple VLANs on the same trunk, they must use 802.1Q tagging (because you can only have one untagged VLAN per interface/trunk).

* Self IPs are then placed on the VLANs to provide BIG-IP presence and routing/ARP functions, but self IPs are not what provides link resiliency-the trunk does.

Why Option D is correct

* You have two physical interfaces and you want resiliency if one fails # put both interfaces into one trunk with LACP enabled.

* You need both external and internal VLANs on those same two links # both VLANs should be configured as tagged on that trunk, so they can coexist on the same aggregated link.

* If either physical interface fails, the trunk remains up via the remaining interface, keeping both VLANs operational.

Why the other options are incorrect

* A: Two VLANs cannot both be untagged on the same trunk/interface. Only one untagged VLAN is possible; additional VLANs must be tagged.

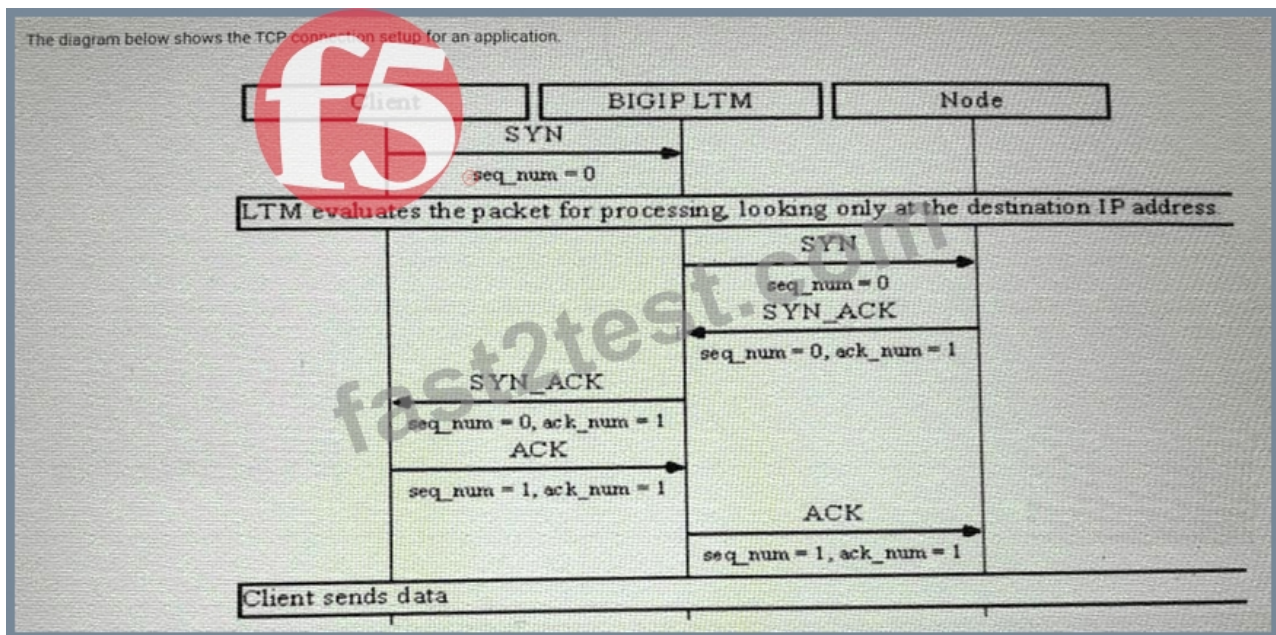
* B: Two trunks "each with one VLAN" would typically mean splitting VLANs across separate trunks.

With only two interfaces total, that becomes one interface per trunk-if one interface goes down, the VLAN on that interface is down (no redundancy for that VLAN).

* C: Same redundancy problem as B, and disabling LACP removes the negotiated aggregation behavior expected when the switch engineer specifically requested LACP.

NEW QUESTION # 11

The diagram below shows the TCP connection setup for an application.



Which of the following virtual server types applies? (Choose one answer)

- A. Standard virtual server
- B. Stateless virtual server
- C. Forwarding IP virtual server

Answer: C

Explanation:

The diagram illustrates a specific TCP handshake sequence where the BIG-IP system acts as a transparent forwarder rather than a full proxy. The key indicators that identify this as a Forwarding (IP) virtual server are as follows:

* **Initial Packet Processing:** The diagram explicitly states that the LTM evaluates the packet looking only at the destination IP address. This is the fundamental characteristic of a Forwarding IP virtual server, which uses the system's routing table to make forwarding decisions instead of load balancing to a pool of members.

* **Handshake Sequence:** Unlike a Standard virtual server, which completes the three-way handshake with the client (SYN, SYN-ACK, ACK) before initiating a separate connection to the server, the Forwarding IP virtual server passes the client's original SYN packet directly to the destination node.

* **Response Timing:** The BIG-IP system waits for the SYN-ACK from the destination node before it sends a SYN-ACK back to the client. It essentially "passes through" the handshake signals while still maintaining a state entry in the connection table to track the flow.

* **Packet-by-Packet Logic:** While it tracks the state, it does not perform address translation (unless SNAT is specifically configured) or deep packet inspection like a full proxy would.

Why other options are incorrect:

* **Standard virtual server:** A Standard virtual server is a "full proxy." It would finish the handshake with the client first and only then open a second, independent TCP connection to the backend server.

* **Stateless virtual server:** A stateless virtual server does not track connections in the connection table.

The diagram shows the system meticulously passing sequence numbers (seq_num) and acknowledgment numbers (ack_num) between the two sides, which requires stateful tracking of the TCP flow.

NEW QUESTION # 12

When using the setup utility to configure a redundant pair, you are asked to provide a "Failover Peer IP". Which address is this?

- A. an address on the current system used to listen for failover messages from the partner BIG-IP
- B. an address of the other system in a redundant pair configuration
- C. an address of the other system in its management network
- D. an address on the current system used to initiate mirroring and network failover heartbeat messages

Answer: B

Explanation:

When establishing a redundant pair, each device must know where to send its health heartbeats and sync data.

* The Peer IP: The Failover Peer IP is the IP address belonging to the other BIG-IP device in the HA pair.

This is typically a self-IP on a dedicated "HA" or "Internal" VLAN, or the Management IP.

* Purpose: It identifies the destination for the "Heartbeat" (the "Are you alive?" check).

* Setup Context: During the initial setup, you tell Device A to look for Device B at its "Failover Peer IP," and you tell Device B to look for Device A at its respective "Failover Peer IP."

NEW QUESTION # 13

Refer to the exhibit above.

Local Traffic » Pools : Pool List » docker_www_farm

Properties Members Statistics

Load Balancing

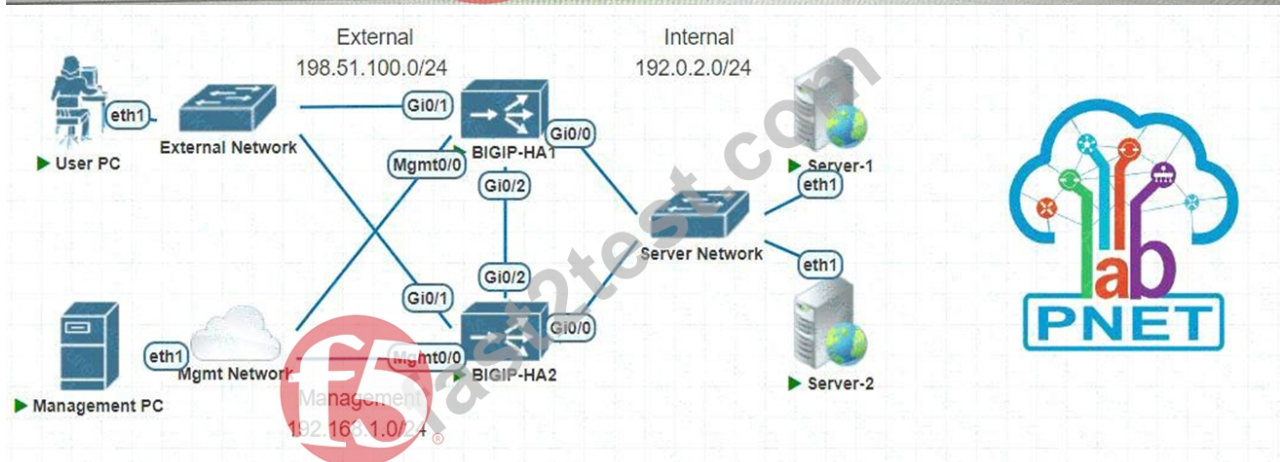
Load Balancing Method: Round Robin

Priority Group Activation: Less than... 2 Available Member(s)




Update

Current Members

✓	▼	Status	Member	Address	Service Port	FQDN	Ephemeral	Ratio	Priority Group	Connection Limit	Partition / Path
<input type="checkbox"/>		●	serv1:80	192.168.30.11	80		No	1	2 (Active)	0	Common
<input type="checkbox"/>		●	serv2:80	192.168.30.12	80		No	1	2 (Inactive)	0	Common
<input type="checkbox"/>		●	serv3:80	192.168.30.13	80		No	1	1 (Active)	0	Common
<input type="checkbox"/>		●	serv4:80	192.168.30.14	80		No	1	1 (Active)	0	Common

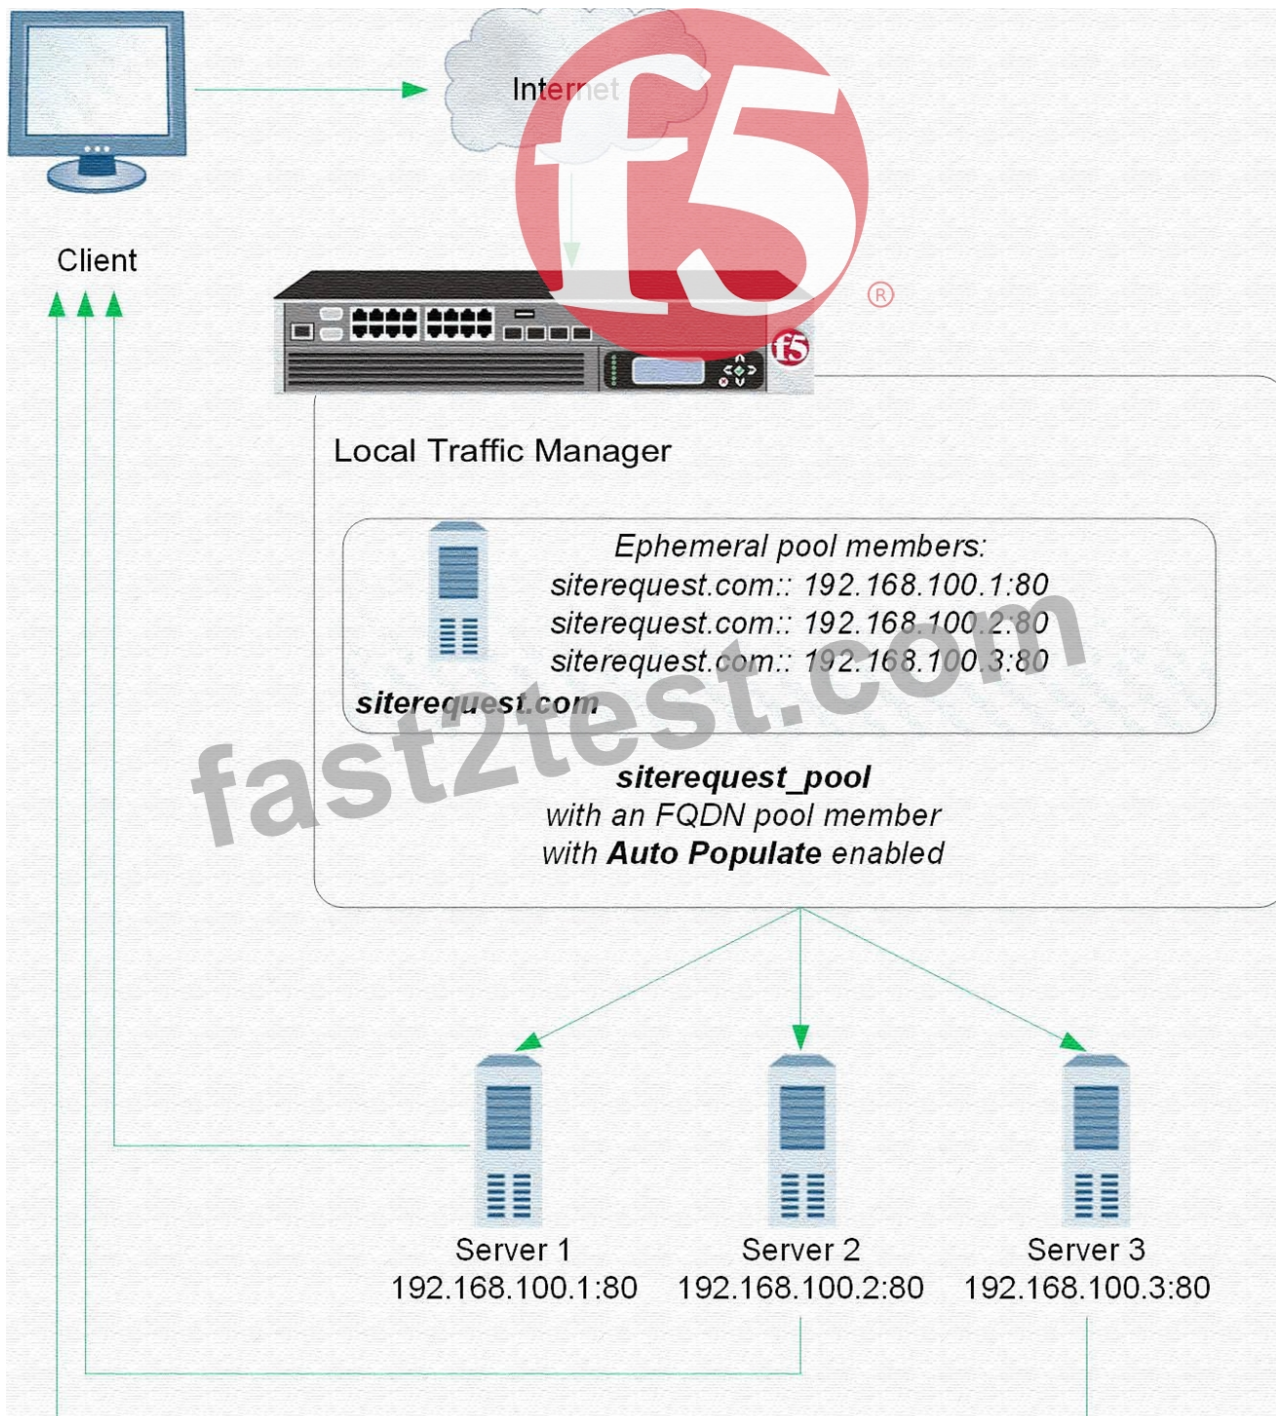


Member Properties

Node Name	10.1.20.11
Address	10.1.20.11
Service Port	80
Partition / Path	Common
Description	<input type="text"/>
Parent Node	 10.1.20.11
Availability	 Available (Enabled) - Pool member is available 2018-05-29 16:56:28
Health Monitors	 http
Monitor Logging	<input type="checkbox"/> Enable
Current Connections	0
State	<input checked="" type="radio"/> Enabled (All traffic allowed) <input type="radio"/> Disabled (Only persistent or active connections allowed) <input type="radio"/> Forced Offline (Only active connections allowed)

Configuration: Basic

Ratio	<input type="text" value="3"/>
Priority Group	<input type="text" value="0"/>
Connection Limit	<input type="text" value="0"/>
Connection Rate Limit	<input type="text" value="0"/>



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 # 14

What command will assist the BIG-IP Administrator in finding the tmm routes when in the TMSH CLI?

- A. list net route
- B. list net
- C. show net route
- D. show net

Answer: C

Explanation:

In the Traffic Management Shell (TMSH), there is a distinct difference between list and show commands.

- * List Command: Used to view the configuration of an object (what is stored in the config file). list net route would show you the static routes you have manually configured.
- * Show Command: Used to view the status, statistics, and real-time state of an object.
- * TMM Routes: Because the routing table (RIB/FIB) is a dynamic entity that includes both static routes and learned routes (such as those from a routing protocol or connected subnets), the show net route command is required to see the active routing table as it

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt,
myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, app.parler.com,
gerardoapp576923.laowaiblog.com, darrenuoee319024.wikiinside.com, Disposable vapes

2026 Latest Fast2test F5CAB2 PDF Dumps and F5CAB2 Exam Engine Free Share: https://drive.google.com/open?id=1qLrdnFMkohWuajhkYp_VnmG5JH-q465L