F5 Networks

F5CAB2 Exam

BIG-IP Administration Data Plane Concepts

Questions & Answers
Demo

Version: 4.0

Question: 1

Which virtual server type is being configured in the screenshot? (Choose one answer.)

A. Standard

B. Forwarding IP

C. Performance Layer 4

Answer: C

Explanation:

Comprehensive and Detailed Explanation (BIG-IP Administration – Data Plane Concepts):

The configuration shown matches a Performance Layer 4 virtual server because it is explicitly using a FastL4 profile:

The screenshot shows Protocol: TCP and Protocol Profile (Client): fastL4.

In BIG-IP data plane terms, FastL4 is the hallmark of a Performance (Layer 4) virtual server, designed to process connections at Layer 4 with minimal overhead (high throughput/low latency) compared to full proxy L7 processing.

The screenshot also shows HTTP Profile (Client): None (and HTTP server profile effectively not in use).

A Standard virtual server commonly uses full-proxy features and frequently includes L7 profiles (like HTTP) when doing HTTP-aware load balancing, header manipulation, cookie persistence, etc. In contrast, a Performance L4 virtual server typically does not use an HTTP profile because it is not doing HTTP-aware (Layer 7) processing.

It is not a Forwarding IP virtual server:

A Forwarding (IP) virtual server is used to route/forward packets (often without load balancing to pool members in the same way as Standard/Performance VS) and is selected by choosing a forwarding type. The presence of a TCP protocol with a FastL4 client profile aligns with a Layer 4 load-balancing style virtual server, not a packet-forwarding virtual server type.

Conclusion: Because the configuration is TCP-based and explicitly uses fastL4 with no HTTP profile, the expected BIG-IP virtual server type is Performance Layer 4 (Option C).

Question: 2

A development team needs to apply a software fix and troubleshoot one of its servers. The BIG-IP Administrator needs to immediately remove all connections from the BIG-IP system to the back-end server. The BIG-IP Administrator checks the virtual server configuration and finds that a persistence profile is assigned to it.

What should the BIG-IP Administrator do to meet this requirement? (Choose one answer)

A. Set the pool member to a Forced Offline state and manually delete existing connections through the command line

B. Set the pool member to an Offline state and manually delete existing connections through the command line

C. Set the pool member to a Forced Offline state

D. Set the pool member to a Disabled state

Answer: C

Explanation:

Comprehensive and Detailed Explanation (BIG-IP Administration – Data Plane Concepts):

In BIG-IP traffic management, persistence profiles cause existing client connections (and subsequent requests) to be repeatedly sent to the same pool member. When persistence is enabled, simply preventing new connections is not sufficient if the requirement is to immediately remove all existing connections.

Key behavior of pool member states:

Forced Offline

Immediately removes the pool member from load balancing.

Terminates all existing connections, regardless of persistence.

Prevents new connections from being established.

This is the correct state when urgent maintenance or troubleshooting is required.

Disabled

Prevents new connections from being sent to the pool member.

Allows existing connections to continue, which is not acceptable when persistence is configured and connections must be cleared immediately.

Offline (non-forced)

Similar to Disabled behavior depending on context.

Does not guarantee immediate termination of existing connections.

Manually deleting connections via the command line

Is unnecessary and operationally inefficient.

BIG-IP already provides a supported mechanism (Forced Offline) to cleanly and immediately remove traffic.

Conclusion:

To immediately remove all existing connections, including those maintained by persistence, the BIG-IP Administrator must set the pool member to a Forced Offline state. This directly satisfies the requirement without additional manual steps.

Question: 3

Refer to the exhibit.

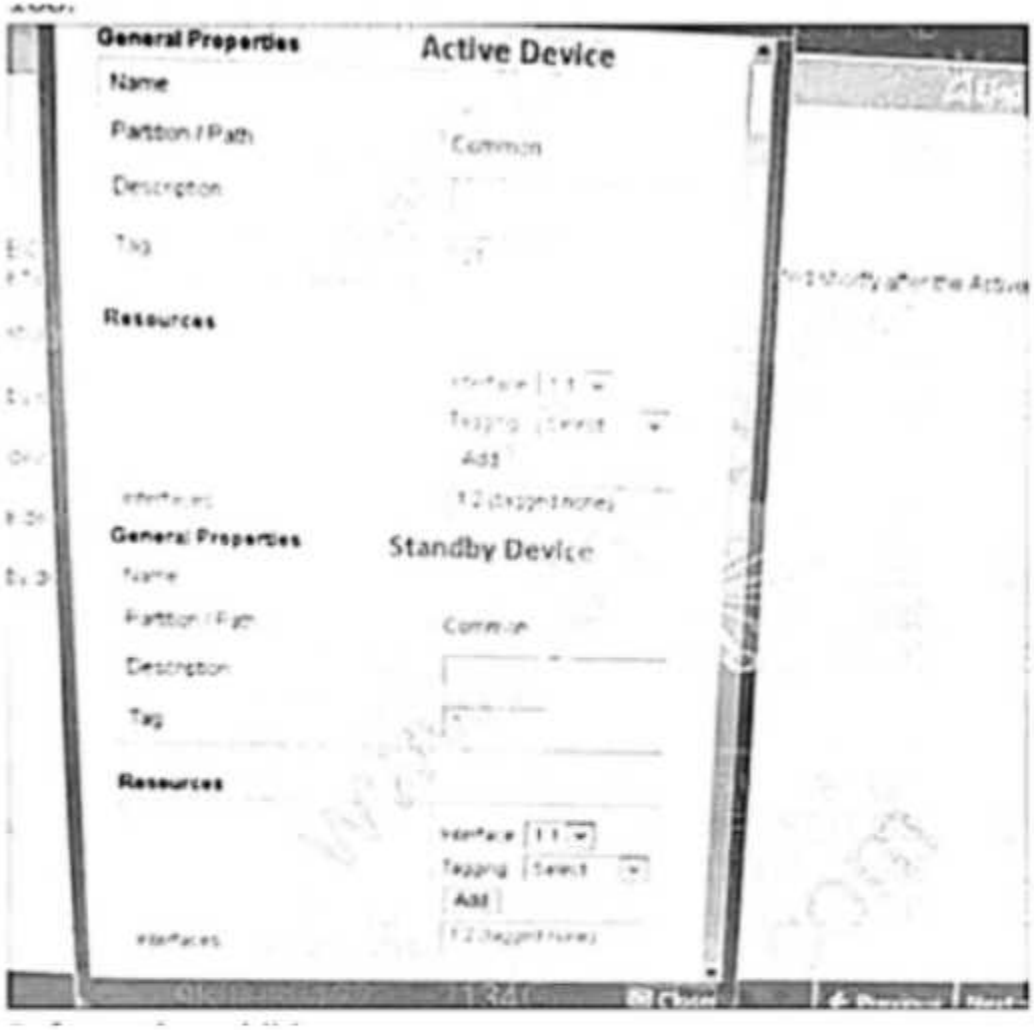
General Properties
Active Device
Name
Partition / Path
Common
Description
Tag
Resources
Add
General Properties
Standby Device
Name
Partition / Path
Common
Description
Tag
Resources
Add
interfaces

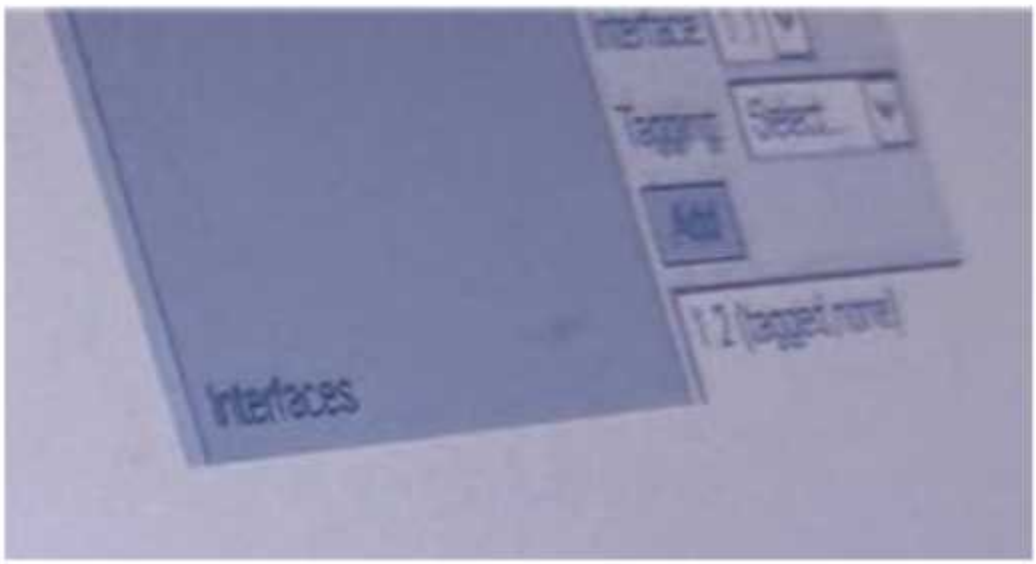
interfaces

During a planned upgrade to a BIG-IP HA pair running Active/Standby, an outage to application traffic is reported shortly after the Active unit is forced to Standby. Reverting the failover resolves the outage. What should the BIG-IP Administrator modify to avoid an outage during the next failover event? (Choose one answer)

A. The Tag value on the Standby device

B. The interface on the Active device to 1.1

C. The Tag value on the Active device

D. The Interface on the Standby device to 1.1

Answer: D

Explanation:

Comprehensive and Detailed Explanation (BIG-IP Administration – Data Plane Concepts):

In an Active/Standby BIG-IP design, application availability during failover depends on both units having equivalent data-plane connectivity for the networks that carry application traffic. Specifically:

VLANs are bound to specific interfaces (and optionally VLAN tags).

Floating self IPs / traffic groups move to the new Active device during failover.

For traffic to continue flowing after failover, the new Active device must have the same VLANs available on the correct interfaces that connect to the upstream/downstream networks.

What the symptom tells you:

Traffic works when Device A is Active

Traffic fails when Device B becomes Active

Failback immediately restores traffic

This pattern strongly indicates the Standby unit does not have the VLAN connected the same way (wrong physical interface assignment), so when it becomes Active, it owns the floating addresses but cannot actually pass traffic on the correct network segment.

Why Interface mismatch is the best match:

If the Active unit is already working, its interface mapping is correct.

The fix is to make the Standby unit's VLAN/interface assignment match the Active unit.

That corresponds to changing the Standby device interface to 1.1.

Why the Tag options are less likely here (given the choices and the exhibit intent):

Tag issues can also break failover traffic, but the question/options are clearly driving toward the classic HA requirement: consistent VLAN-to-interface mapping on both devices so the data plane remains functional after the traffic group moves.

Conclusion: To avoid an outage on the next failover, the BIG-IP Administrator must ensure the Standby device uses the same interface (1.1) for the relevant VLAN(s) that carry the application traffic, so when it becomes Active it can forward/receive traffic normally.

Question: 4

Active connections to pool members are unevenly distributed. The load balancing method is Least Connections (member). Priority Group Activation is disabled.

What is a potential cause of the uneven distribution? (Choose one answer)

A. Priority Group Activation is disabled

B. SSL Profile Server is applied

C. A persistence profile is applied

D. Incorrect load balancing method

Answer: C

Explanation:

Comprehensive and Detailed Explanation (BIG-IP Administration – Data Plane Concepts):

With Least Connections (member), BIG-IP attempts to send new connections to the pool member with the fewest current connections. In a perfectly "stateless" scenario (no affinity), this often trends toward a fairly even distribution over time.

However, persistence overrides load balancing:

When a persistence profile is applied, BIG-IP will continue sending a client (or client group) to the same pool member based on the persistence record (cookie / source address / SSL session ID, etc.).

This means even if another pool member has fewer connections, BIG-IP may still select the persisted member to honor session affinity.

The result can be uneven active connection counts, even though the configured load balancing method is Least Connections.

Why the other options are not the best cause:

A . Priority Group Activation is disabled

Priority Group Activation only affects selection when priority groups are configured; disabling it does not inherently create uneven distribution under Least Connections.

B . SSL Profile Server is applied

A server-side SSL profile affects encryption to pool members, but it does not by itself cause skewed selection across pool members. (Skew could happen indirectly if members have different performance/latency, but that's not the primary, expected exam answer.)

D . Incorrect load balancing method

Least Connections is a valid method and does not itself explain unevenness unless something is overriding it (like persistence) or pool members are not all eligible.

Conclusion:

A persistence profile is the most common and expected reason that active connections become unevenly distributed, because persistence takes precedence over the Least Connections load-balancing decision.

Question: 5

and their status/statistics]

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. Increase the MTU on the VLAN using interface 1.1

B. Create a trunk object with two interfaces

C. Assign two interfaces to the VLAN

D. Set the media speed of interface 1.1 manually

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.