



Vendor: F5

Exam Code: F5CAB3

Exam Name: BIG-IP Administration Data Plane Configuration

Version: DEMO

QUESTION 1

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. Disable Address Translation
- 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. Set Address Translation to SNAT and have a self-IP configured in the same subnet as the servers

Answer: D

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.

QUESTION 2

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

Answer: D

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.

QUESTION 3

A BIG-IP Administrator configures a node with a standard icmp Health Monitor. The Node shows as DOWN although the Backend Server is configured to answer ICMP requests. Which step should the administrator take next to find the root cause of this issue?

- A. Run a qkview
- B. Run a tcpdump
- C. Run an ssldump
- D. Run a curl

Answer: B

Explanation:

In the F5 BIG-IP ecosystem, a standard ICMP health monitor functions by sending an ICMP echo request to a target node and expecting an ICMP echo reply within a specified timeout period. When a node is marked "DOWN" despite the backend server being configured to respond to ICMP, the issue typically lies in the network path or the specific packet exchange between the BIG-IP's self IP and the node's IP. Running a tcpdump is the most effective next step because it provides a real-time packet capture of the actual monitor traffic leaving the BIG-IP and any return traffic coming back from the server. This allows the administrator to verify if the BIG-IP is actually sending the echo request, if the request is reaching the server, and if the server is indeed replying or if the reply is being dropped by an intermediate firewall or a security policy.

QUESTION 4

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?

General Properties	
Name	HTTPS_Virtual
Partition / Path	Common
Description	
Type	Standard
Source Address	0.0.0.0/0
Destination Address/Mask	10.1.1.4
Service Port	443 HTTPS
Notify Status to Virtual Address	☒
Availability	● Available (Enabled) - The virtual server is available
Synccookie Status	Off
State	Enabled

Configuration: Basic

Configuration: Basic	
Protocol	TCP
Protocol Profile (Client)	tcp
Protocol Profile (Server)	(Use Client Profile)
HTTP Profile	http
FTP Profile	None
RTSP Profile	None
SSH Proxy Profile	None
SSL Profile (Client)	Selected Available << >> / Common clientssl clientssl-insecure-compatible clientssl-secure crypto-server-default-clientssl

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

Answer: A

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.

QUESTION 5

A BIG-IP Administrator creates an HTTP Virtual Server using an iApp template. After the Virtual Server is created, the user requests to change the destination IP addresses. The BIG-IP Administrator tries to change the destination IP address from 10.1.1.1 to 10.2.1.1 in Virtual Server settings, but receives the following error: "The application service must be updated using an application management interface." What is causing this error?

- A. The Application Services have Strict Updates enabled.
- B. The IP addresses are already in use.
- C. The IP addresses used are NOT from the same subnet as the Self IP.
- D. The Application Service was NOT deleted before making the IP address change.

Answer: A

Explanation:

In F5 BIG-IP administration, iApps are designed to manage complex application configurations as a single unit. When an iApp is deployed, it creates an "Application Service" object that owns all the associated LTM objects, such as Virtual Servers, Pools, and Nodes. By default, these iApps are created with Strict Updates enabled. Strict Updates is a safety mechanism that prevents administrators from making manual "out-of-band" changes to the individual objects created by the iApp. The system enforces this because manual changes would be overwritten the next time the iApp template is updated or re-entered.

When the administrator attempts to change the destination IP address directly on the Virtual Server object, the BIG-IP system checks the "Strict Updates" flag. If it is set to "Enabled," the system blocks the modification and generates the error message stating the service must be updated via the application management interface. To resolve this, the administrator must navigate to the iApp >> Application Services menu, select the specific application service, and go to the "Reconfigure" tab. Within the iApp configuration form, the destination IP can be safely changed. Alternatively, if the administrator specifically wants to manage the objects manually and forgo the benefits of the iApp template management, they could disable "Strict Updates" in the iApp properties, though this is generally discouraged as it breaks the template's logic. The error is not related to subnetting or duplicate IPs, but strictly to the configuration authority assigned to the iApp service.

QUESTION 6

Refer to the exhibit. A BIG-IP Administrator configures a new VLAN on an HA pair of devices that does NOT yet have any traffic. This action causes the assigned traffic group to fail over to the standby device. Which VLAN setting should be changed to prevent this issue?

The screenshot displays a network configuration interface for a VLAN. The top navigation bar shows 'Network > VLANs > VLAN List > External'. Below this, there are tabs for 'Properties' (selected) and 'Layer 2 Static Forwarding Table'. The 'General Properties' section includes fields for Name (External), Partition / Path (Common), Description (empty), and Tag (4093). The 'Resources' section shows a list of interfaces with '1.2 (untagged)' selected. The 'Advanced' configuration section is expanded, showing various settings: Source Check (unchecked), MTU (1500), MAC Address (00:50:56:9d:3c:47), Fail-safe (checked), Fail-safe Timeout (90), Action (Reboot), Auto Last Hop (Default), CMP Hash (Default), DAG Tunnel (Outer), IPv6 Prefix Length (128), and DAG Adjustment (None).

- A. Auto Last Hop
- B. Source Check
- C. Fail-safe
- D. Customer Tag

Answer: C

Explanation:

The exhibit shows the advanced configuration of a VLAN where the Fail-safe option is checked. VLAN Fail-safe is a high-availability feature used to monitor network connectivity on a specific VLAN. When enabled, the BIG-IP system monitors the VLAN for network traffic. If the system does not detect any "useful" traffic on the VLAN within the specified Fail-safe Timeout (which is 90 seconds in the exhibit), it attempts to generate traffic by pinging the default gateway or other devices. If it still detects no traffic, the BIG-IP concludes that the VLAN is unreachable or the network interface has failed, and it triggers a "Fail-safe Action"--in this case, "Reboot" or a failover to the peer device in the HA group.

Because the administrator has just created a new VLAN that "does NOT yet have any traffic," the Fail-safe mechanism triggers immediately after the 90-second timeout period. Since no devices are yet communicating on this VLAN, the BIG-IP incorrectly assumes there is a hardware or cabling failure and forces a failover to ensure the standby device (which might have better connectivity) takes over. To prevent this unwanted failover, the administrator should uncheck the Fail-safe box for that specific VLAN until the VLAN is fully populated with active nodes and

regular traffic. Once the application is live and traffic is flowing, Fail-safe can be re-enabled to provide an additional layer of redundancy. Auto Last Hop (Option A) and Source Check (Option B) are routing and security features that do not trigger HA failover events.

QUESTION 7

In a pool there are 2 pool members out of the 5 members that are older servers. The number of connections these can handle is less than the other 3 pool members. Which load balancing method would allow more traffic to be directed to the newer servers?

- A. Weighted Least Connections (member)
- B. Global Availability
- C. Round Robin
- D. Least Connections (member)

Answer: A

Explanation:

When dealing with heterogeneous server hardware where some servers are more powerful than others, a dynamic load balancing method that accounts for both current load and server capacity is required. The Weighted Least Connections (member) method is the most appropriate choice. This method works by tracking the number of active connections to each pool member and then "weighting" that number based on a user-defined Ratio value assigned to the member. For example, the administrator can assign a higher Ratio to the three newer, more powerful servers and a lower Ratio to the two older servers. The BIG-IP then uses a formula to calculate which server should receive the next connection, ensuring that the newer servers handle a proportionately larger share of the total concurrent connections.

Standard Round Robin (Option C) would be ineffective because it distributes connections strictly sequentially (1, 2, 3, 4, 5) without regard for the servers' capacity or current load, which would eventually overwhelm the older servers. Least Connections (member) (Option D) is better than Round Robin because it picks the server with the fewest active connections, but it still assumes all servers are equal; it would try to keep the connection counts identical across all 5 servers, which would still stress the older hardware more than the new. Global Availability (Option B) is a GSLB (DNS-based) method used for multi-site redundancy, not for local pool member load balancing. By using Weighted Least Connections, the administrator achieves a balance where the more capable servers take the brunt of the work while the older servers are utilized only to their specific safe capacity.

QUESTION 8

All pool members are online. All other virtual server settings are at default. What might alter the load balancing behavior?

- A. Adding a oneconnect profile
- B. Adding a persistence profile
- C. Enabling a fallback host in the http profile
- D. Enabling SNAT automap

Answer: B

Explanation:

In a default BIG-IP configuration, the system utilizes the Load Balancing Method (typically Round Robin) to distribute each new connection across available pool members. However, the introduction of a persistence profile fundamentally changes this behavior. Persistence (also known as "stickiness") ensures that once a client has been load balanced to a specific pool member, all subsequent requests from that same client during a defined session or timeout

period are directed to that same member, bypassing the standard load balancing algorithm. This is critical for applications that maintain state, such as shopping carts or authenticated sessions, where moving a user to a different server would result in a loss of session data.

QUESTION 9

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

Answer: BD

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.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



Microsoft



ORACLE



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14