F5.F5CAB3.v2026-07-16.q48

Exam Code:	F5CAB3
Exam Name:	BIG-IP Administration Data Plane Configuration
Certification Provider:	F5
Free Question Number:	48
Version:	v2026-07-16
# of views:	126
# of Questions views:	480
https://www.exam-tests.com/F5CAB3-exam/F5.F5CAB3.v2026-07-16.q48.html	

NEW QUESTION: 1
A configuration change is made on the standby member of a device group. What is displayed as "Recommended Action" on the Device Management Overview screen?
A. Force active member of device group to standby
B. Activate device with the most recent configuration
C. Synchronize the active member configuration to the group
D. Synchronize the standby member configuration to the group
Answer: D (LEAVE A REPLY)
The BIG-IP system uses a centralized management framework to ensure that all devices within a Sync-Failover group share a consistent configuration. When an administrator makes a change on any member of the group-whether it is the active or the standby device-the system detects a "ConfigSync" mismatch. The "Device Management >> Overview" screen tracks these changes by comparing the commit ID and timestamps of the configurations across all peers.
If a change is made on the standby member, that device now possesses a more recent configuration than the other members of the group. Consequently, the BIG-IP GUI will display a status of "Changes Pending" and suggest a Recommended Action to resolve the discrepancy. In this scenario, the correct action is to Synchronize the standby member configuration to the group. This push operation will copy the updated configuration from the standby device to the active device (and any other peers), bringing the entire cluster back into a "In Sync" status. It is important to note that BIG-IP allows bi-directional synchronization; you do not have to be on the active device to push a configuration. However, administrators must be cautious: choosing Option C (Synchronizing the active member to the group) would overwrite the changes just made on the standby device with the older configuration from the active device, effectively reverting the changes. The Recommended Action always points toward the direction that propagates the most recent change to the rest of the group.

NEW QUESTION: 2

A BIG-IP Administrator needs to apply persistence to a virtual server that is configured as a Performance (Layer 4) virtual server that allows access to a secure (TLS) e-commerce website.
What type of persistence profile can be used? (Choose one answer)
A. Source Address Affinity
B. Cookie persistence
C. Microsoft RDP persistence
D. Host persistence
Answer: A (LEAVE A REPLY)
A Performance (Layer 4) virtual server does not inspect or process application-layer data such as HTTP headers or cookies. Therefore, only Layer 4-compatible persistence methods can be used.
According to the BIG-IP Administration: Data Plane Configuration documentation:
Source Address Affinity persistence operates at Layer 4 and uses the client IP address to maintain session persistence.
It is fully compatible with Performance (Layer 4) virtual servers.
It works regardless of encryption, making it suitable for TLS-secured applications.
Why the other options are incorrect:
B . Cookie persistence
Requires an HTTP profile and Layer 7 inspection, which is not supported on Performance virtual servers.
C . Microsoft RDP persistence
Is protocol-specific and not applicable to web-based TLS traffic.
D . Host persistence
Requires HTTP host header inspection, which is not available at Layer 4.
Correct Resolution:
Source Address Affinity persistence is the appropriate choice for maintaining persistence on a Performance (Layer 4) virtual server handling TLS traffic.
Below is Batch 1 (Questions 1-10) extracted only from your uploaded document that are directly related to BIG-IP Administration: Data Plane Configuration topics (Virtual Servers, Pools, Load Balancing, Monitors, Persistence, SNAT, Profiles).
I have excluded system-only, licensing, support, hardware, HA management-only, and admin UI questions that are not Data Plane-focused.
Source: Your uploaded TMOS Administration v2.0 document
✅ BATCH 1 (10 Questions)

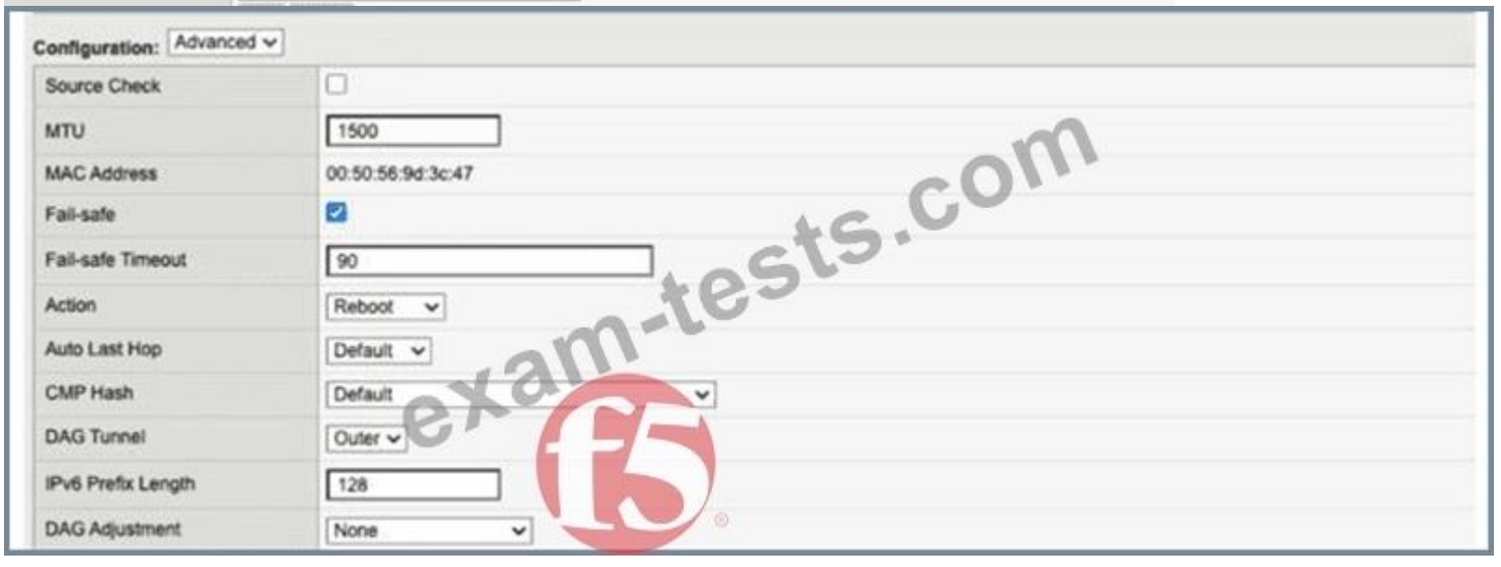
NEW QUESTION: 3
A Standard Virtual Server for a web application is configured with Automap for Source Address Translation.
The original client IP must be known by backend servers.
What should the BIG-IP Administrator configure?
A. Performance (HTTP) Virtual Server
B. HTTP profile to insert X-Forwarded-For
C. HTTP Transparent profile
D. SNAT pool using client IP

Answer: B (LEAVE A REPLY)
The X-Forwarded-For header preserves the original client IP when SNAT is enabled.

NEW QUESTION: 4
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. ICMP + TCP monitor
C. HTTP monitor with custom send/receive
D. TCP monitor
Answer: C (LEAVE A REPLY)
Application-level monitors ensure that only the affected service is marked down, not the entire node.

NEW QUESTION: 5
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?

A. Auto Last Hop

B. Source Check

C. Fail-safe

D. Customer Tag

Answer: C (LEAVE A REPLY)

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.

NEW QUESTION: 6

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.

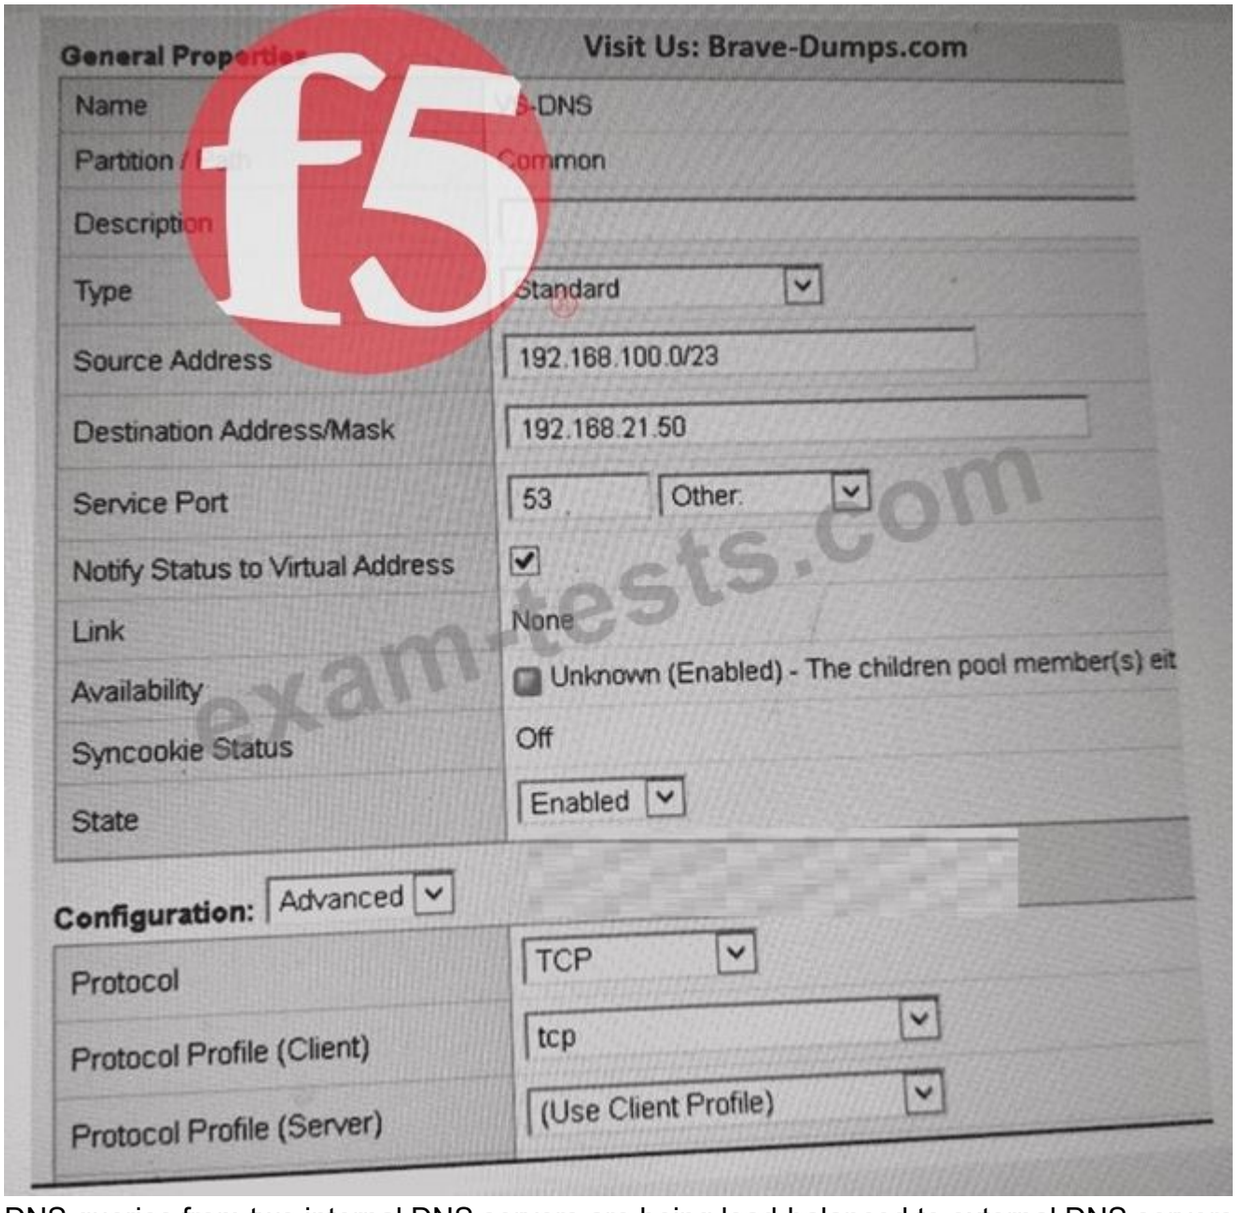
Answer: B,D (LEAVE A REPLY)

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.

NEW QUESTION: 7
Refer to the exhibit.



DNS queries from two internal DNS servers are being load-balanced to external DNS servers via a virtual server on a BIG-IP device. The DNS queries originate from:
192.168.10.100
192.168.10.200
and target:
192.168.2.150

All DNS queries destined for the external DNS servers fail.
Which property change should the BIG-IP Administrator make in the Virtual Server to resolve this issue? (Choose one answer)
A. Protocol profile (Client) to DNS_OPTIMIZED
B. Type to Performance (HTTP)
C. Source Address to 192.168.10.0/24
D. Protocol to UDP
Answer: D (LEAVE A REPLY)
DNS traffic is primarily transported using UDP port 53. In the exhibit, the Virtual Server is configured with the Protocol set to TCP, which prevents standard DNS queries from being processed correctly. BIG-IP Virtual Servers must be configured with the correct Layer 4 protocol to match the application traffic they are handling.
According to the BIG-IP Administration: Data Plane Configuration documentation:
The Protocol setting on a Virtual Server defines whether traffic is processed as TCP, UDP, or another supported transport protocol.
Standard DNS queries and responses use UDP, while TCP is only required for DNS zone transfers (AXFR) or exceptionally large responses.
When a DNS Virtual Server is incorrectly configured with TCP, UDP-based DNS queries are dropped, causing all requests to fail.
Why the other options are incorrect:
A). Protocol profile (Client) to DNS_OPTIMIZEDA DNS profile enhances DNS functionality but does not correct an incorrect transport protocol configuration.
B). Type to Performance (HTTP)Performance (HTTP) Virtual Servers are designed for HTTP traffic and are not suitable for DNS services.
C). Source Address to 192.168.10.0/24The existing source IPs already fall within the allowed range, so this setting does not address the failure.
Correct Resolution:
Changing the Protocol to UDP aligns the Virtual Server with standard DNS transport requirements, allowing DNS queries to be successfully processed and load-balanced.

NEW QUESTION: 8
A BIG-IP Administrator needs to modify a virtual server that will offload web traffic compression tasks from the target server. Which two profiles must the BIG-IP Administrator apply to a virtual server to enable compression? (Pick the 2 correct responses below)
A. Persistence profile
B. Compression profile
C. Server SSL profile
D. Stream profile
E. HTTP profile
Answer: B,E (LEAVE A REPLY)
HTTP Compression on a BIG-IP system is an optimization feature that reduces the size of HTTP responses sent to clients, thereby saving bandwidth and improving page load times for end-users. To implement this

functionality on a Virtual Server, the system requires two specific profiles to work in tandem: an HTTP profile and a Compression profile.
The HTTP profile (Option E) is a prerequisite because compression is an application-layer (Layer 7) operation. The BIG-IP must be "HTTP-aware" to parse the incoming requests and outgoing responses. The HTTP profile allows the BIG-IP to inspect the Accept-Encoding header from the client (to see if the client supports gzip or deflate) and to modify the Content-Encoding and Vary headers in the server's response.
Without the HTTP profile, the Virtual Server operates at Layer 4 (TCP) and sees the traffic only as an opaque stream of bytes, making compression impossible.
The Compression profile (Option B) contains the actual settings and engine for the task. It defines which content types should be compressed (e.g., text/html, application/javascript), the compression level, and the minimum content length required before compression is applied. While a Stream profile (Option D) is used for data replacement (search and replace) and SSL profiles (Option C) are used for encryption, they do not provide compression logic. By applying both the HTTP and Compression profiles, the BIG-IP can successfully intercept server responses, compress the payload according to the configured policy, and deliver the optimized data to the user, effectively offloading the CPU-intensive compression work from the backend web servers.

NEW QUESTION: 9
Exhibit:



Due to a change in application requirements, a BIG-IP Administrator needs to modify the configuration of a Virtual Server to include a Fallback Persistence Profile. Which persistence profile type should the BIG-IP Administrator use for this purpose?
A. Universal
B. Source Address Affinity
C. SSL
D. Hash
Answer: (SHOW ANSWER)
In a BIG-IP environment, a Fallback Persistence Profile is utilized as a secondary "stickiness" mechanism when the primary (Default) persistence method fails to provide a valid persistence record. For example, if a Virtual Server uses HTTP Cookie Persistence as its primary method, but a client's browser has cookies disabled, the BIG-IP will be unable to find a persistence cookie in the request. Without a fallback method, the system would treat every request from that client as a new, independent connection, potentially breaking the application session.

Source Address Affinity (also known as Source Address Persistence) is the most common and standard choice for a fallback profile. It operates at the network layer (Layer 3) by tracking the client's source IP address. Because every IP packet contains a source address, this method is virtually guaranteed to work even when application-layer data (like Cookies or SSL Session IDs) is missing or encrypted beyond the BIG-IP's visibility. While Universal (Option A) and Hash (Option D) profiles are highly flexible and can use iRules to persist on almost any data, they require specific configuration and logic that may not always be present or valid. SSL persistence (Option C) relies on the SSL Session ID, which frequently changes due to modern browser security practices (session renegotiation), making it less reliable than Source Address Affinity. By configuring Source Address Affinity as the fallback, the administrator ensures that the BIG-IP has a "safety net" to maintain session integrity based on the client's IP address when the more granular cookie-based persistence is unavailable.

NEW QUESTION: 10

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 (LEAVE A REPLY)

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.

Why other options are incorrect:

A: Disabling address translation would ensure the server-side traffic uses the client IP, making asymmetric routing inevitable in this scenario.

B: This is technically contradictory; " Auto Map " specifically uses existing self-IPs and does not require or use a " SNAT pool " configuration.
C: While using a specific translation address can work, it does not inherently guarantee the Layer 2/Layer 3 reachability mentioned in the scenario as effectively as ensuring the self-IP is correctly placed in the server ' s subnet.

NEW QUESTION: 11
A set of servers is used for an FTP application as well as an HTTP website via separate BIG-IP Pools. The server support team reports that some servers are receiving a lot more traffic than others. Which Load Balancing Method should the BIG-IP Administrator apply to even out the connection count?
A. Ratio (Node)
B. Least Connections (Member)
C. Ratio (Member)
D. Least Connections (Node)
Answer: D (LEAVE A REPLY)
Similar to the logic required for managing multi-service backend environments, the issue described-where servers hosting multiple protocols like FTP and HTTP are experiencing uneven distribution-stems from the BIG-IP's default behavior of treating each pool independently. If the administrator uses a member-based load balancing method, the BIG-IP distributes HTTP traffic regardless of how much FTP traffic that same physical server is currently processing.
To resolve this, the administrator must utilize the Least Connections (Node) method. By switching both the HTTP and FTP pools to this algorithm, the BIG-IP begins to make load balancing decisions based on the total combined connection count for the IP address of each server. When a new HTTP request arrives, the BIG-IP checks which server has the fewest total connections (including existing FTP sessions). This prevents a server that is already busy with long-lived FTP transfers from being overwhelmed by a sudden burst of HTTP requests.
Ratio methods (Options A and C) are static and rely on the administrator manually assigning weights to servers based on their perceived capacity; they do not adapt to real-time fluctuations in traffic volume across different pools. Least Connections (Member) (Option B) remains blind to the "cross-pool" traffic on the same hardware. Only the Node-based Least Connections approach provides the global visibility necessary to "even out" the total resource utilization across servers supporting multiple distinct applications.

NEW QUESTION: 12
A Standard Virtual Server reports poor network performance for Internet-based clients.
What configuration should be applied?
A. Client TCP: f5-tcp-wan / Server TCP: f5-tcp-lan
B. Client TCP: f5-tcp-lan
C. Client TCP: f5-tcp-lan / Server TCP: f5-tcp-wan
D. Client TCP: f5-tcp-optimized
Answer: (SHOW ANSWER)

WAN TCP profiles are optimized for high latency and packet loss typical of Internet clients, while LAN profiles are ideal for backend servers.

NEW QUESTION: 13
Where in the configuration utility should the BIG-IP Administrator verify the pool member currently assigned to a pool is on port 80?
A. Local Traffic > Nodes: Node List. Select the node in question, view the Health Monitor next to Configuration.
B. Local Traffic > Pools: Pool List. Select the pool in question, select Members tab, view the configured Health Monitor.
C. Local Traffic > Pools: Pool List. Select the pool in question, select the Members tab, view the configured Service Port.
Answer: (SHOW ANSWER)
The BIG-IP Configuration Utility (GUI) organizes information hierarchically to allow for granular management of application objects. A Pool is a collection of backend servers (pool members) that provide the same service. To verify the specific network parameters-such as the IP address and the service port-of the servers within a pool, the administrator must navigate to the specific pool's configuration.
The standard procedural path to verify this is Local Traffic > Pools: Pool List, where the administrator selects the specific pool name. Once inside the pool's configuration, the Members tab displays a list of all IP addresses and service ports associated with that pool. Under the "Service Port" column, the administrator can confirm if the member is listening on port 80 (HTTP).
Options A and B are incorrect for this specific verification task. While Nodes (Option A) show the health of a physical server, a node represents only an IP address and does not have a "Service Port" associated with it until it is defined as a pool member. Verifying the Health Monitor (Option B) would tell the administrator how the system is checking the member's status, but it does not definitively show the port on which the member is actually receiving application traffic. In a BIG-IP environment, a pool member is uniquely identified by the combination of its Node IP and its Service Port, and the Members tab is the primary interface for managing and auditing these specific member attributes.

NEW QUESTION: 14
A BIG-IP Administrator is configuring an SSH Pool with five members. Which Health Monitor should be applied to ensure that available pool members are monitored accordingly?
A. UDP
B. HTTP
C. HTTPS
D. TCP
Answer: D (LEAVE A REPLY)
SSH (Secure Shell) operates over TCP port 22 , making the TCP health monitor the appropriate and correct choice for monitoring pool members in an SSH pool. The TCP monitor establishes a basic Layer 4 connection to each pool member on the designated port, verifying that the server is reachable and actively accepting TCP connections. A successful three-way handshake confirms the pool member is available; a failure marks it down and removes it from load balancing rotation.

The remaining monitor types are inappropriate for this use case:
* UDP monitor operates at Layer 4 but uses the connectionless UDP protocol, which is incompatible with SSH since SSH exclusively uses TCP as its transport.
* HTTP monitor sends a Layer 7 GET request expecting an HTTP response - irrelevant for an SSH service that does not speak the HTTP protocol.
* HTTPS monitor functions identically to HTTP but over TLS - again, completely misaligned with SSH traffic and port behaviour.
For application-specific SSH validation beyond basic TCP connectivity, an external monitor or custom monitor could be used; however, TCP remains the standard and correct built-in monitor for SSH pool verification within the BIG-IP platform.
Reference: BIG-IP Administration - Data Plane Configuration, Module: Health Monitors - Monitor Types and Protocol Alignment.

NEW QUESTION: 15
All pool members are online. All other virtual server settings are at default.
What might alter the load balancing behavior? (Choose one answer)
A. Adding a OneConnect profile
B. Enabling SNAT automap
C. Enabling a fallback host in the HTTP profile
D. Adding a persistence profile
Answer: (SHOW ANSWER)
By default, BIG-IP load balancing algorithms (such as Round Robin) distribute connections evenly across all available pool members. However, persistence profiles override normal load balancing decisions by forcing subsequent connections from a client to be sent to the same pool member.
According to the BIG-IP Administration: Data Plane Configuration documentation:
Persistence creates a client-to-server mapping that is honored before load balancing algorithms are applied.
When persistence is enabled, BIG-IP may repeatedly select the same pool member even if others are available.
This directly alters load balancing behavior.
Why the other options are incorrect:
A). Adding a OneConnect profileOneConnect optimizes server-side TCP connections but does not change which pool member is selected.
B). Enabling SNAT automapSNAT affects source address translation, not pool member selection.
C). Enabling a fallback host in the HTTP profileA fallback host is only used when no pool members are available.
Correct Resolution:
Adding a persistence profile alters load balancing behavior by maintaining client affinity to a specific pool member.

NEW QUESTION: 16

A Virtual Server uses an iRule to send traffic to pool members depending on the URI. The BIG-IP Administrator needs to modify the pool member in the iRule.
Which event declaration does the BIG-IP Administrator need to change to accomplish this?
A. CLIENT_ACCEPTED
B. HTTP_RESPONSE
C. HTTP_REQUEST
D. SERVER_CONNECTED
Answer: C (LEAVE A REPLY)
URI-based traffic steering requires inspection of the HTTP request. BIG-IP processes HTTP headers and URIs in the HTTP_REQUEST event. Pool member selection based on URI must occur before the request is sent to the server, making HTTP_REQUEST the correct event.

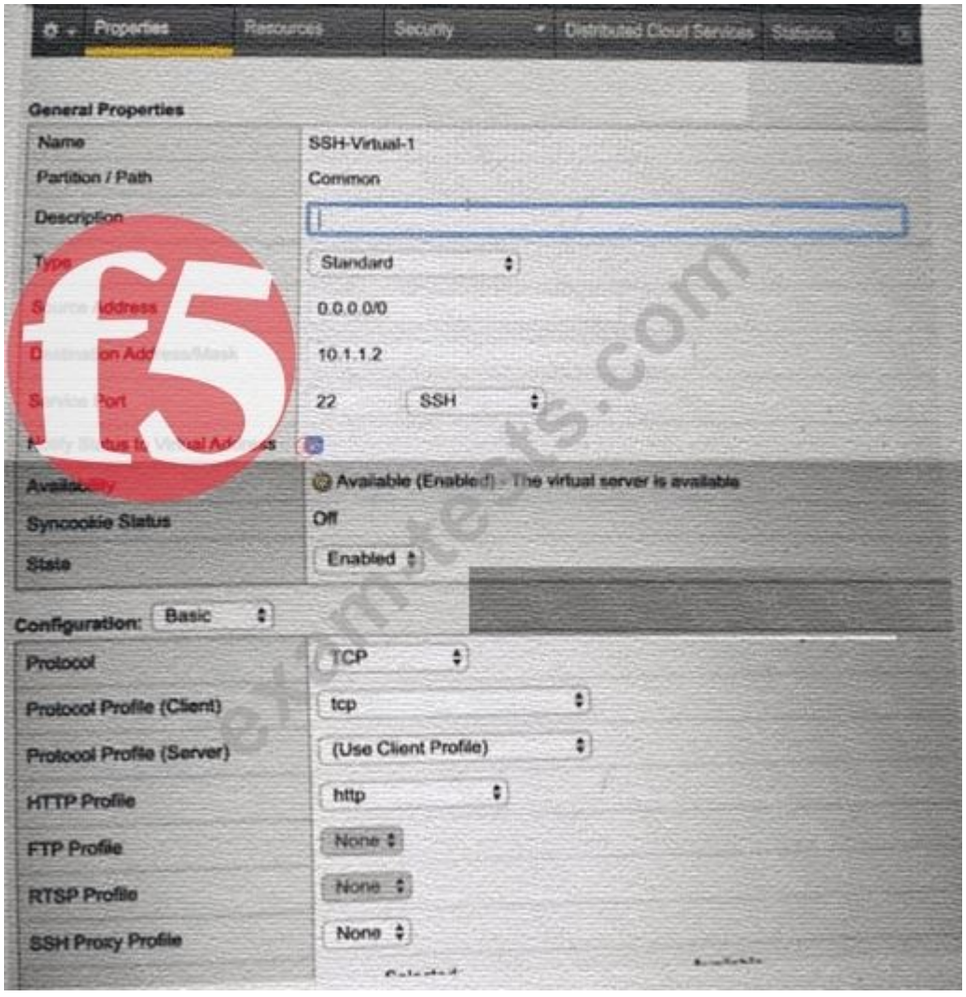
Valid F5CAB3 Dumps shared by BraindumpsPass.com for Helping Passing F5CAB3 Exam! BraindumpsPass.com now offer the **newest F5CAB3 exam dumps**, the BraindumpsPass.com F5CAB3 exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com F5CAB3 dumps with Test Engine here: https://www.braindumpspass.com/F5/F5CAB3-practice-exam-dumps.html (**84** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 17
A Standard Virtual Server for a web application is configured with Automap for the Source Address Translation option. The original source address of the client must be known by the backend servers. What should the BIG-IP Administrator configure to meet this requirement?
A. The Virtual Server type as Performance (HTTP)
B. An HTTP profile to insert the X-Forward-For header
C. A SNAT Pool with the client IP
D. An HTTP Transparent profile
Answer: B (LEAVE A REPLY)
SNAT Automap is a common configuration that replaces the client's original source IP address with one of the BIG-IP's self IP addresses. This ensures that the backend servers send return traffic back through the BIG- IP, which is necessary for the ADC to process the traffic correctly. However, a side effect of SNAT is that the backend servers only see the BIG-IP's IP in their logs, losing visibility into the true identity of the client.
To resolve this while still using SNAT for routing purposes, the administrator must configure the BIG-IP to "pass" the client's IP address at the application layer. This is achieved by using an HTTP Profile with the Insert X-Forwarded-For setting enabled. When this profile is applied to the Virtual Server, the BIG-IP intercepts the HTTP request, adds a header (X-Forwarded-For) containing the client's original IP, and then forwards the modified request to the server. The backend web server can then be configured to read this header and log the original client IP instead of the BIG-IP's SNAT address.

Other options are incorrect for this requirement. Performance (HTTP) (Option A) is a virtual server type optimized for speed but often lacks the full Layer 7 header manipulation capabilities of a Standard Virtual Server. SNAT Pool with the client IP (Option C) is technically impossible as SNAT pools use static, pre- defined IPs. There is no such thing as an HTTP Transparent profile (Option D) in standard BIG-IP administration for this purpose. The X-Forwarded-For header insertion within the HTTP profile is the standard procedural method for maintaining client visibility in SNAT-enabled environments.

NEW QUESTION: 18
A BIG-IP Administrator needs to apply a health monitor for a pool of database servers named DB_Pool that uses TCP port 1521. Where should the BIG-IP Administrator apply this monitor?
A. Local Traffic > Profiles > Protocol > TCP
B. Local Traffic > Pools > DB.Pool > Properties
C. Local Traffic > Nodes > Default Monitor
D. Local Traffic > Pools > DB.Pool > Members
Answer: B (LEAVE A REPLY)
In BIG-IP configuration, health monitors can be applied at three distinct levels: the node, the pool, or the individual pool member. To ensure that a specific application service-in this case, a database service on port 1521-is functioning correctly for the entire pool, the administrator should apply the monitor at the pool level. Navigating to Local Traffic > Pools > DB.Pool > Properties allows the administrator to select one or more monitors from the "Available" list and move them to the "Active" list.
Applying a monitor at the pool property level ensures that the BIG-IP checks the health of every member assigned to that pool using the same logic. If a database-specific monitor (such as a TCP handshake or an Oracle/SQL check) fails for a specific member, the BIG-IP marks that member as "offline" for that specific pool, preventing new connections from being sent to it. While monitors can be applied to Pool Members (Option D) to give different members unique monitoring logic, it is more administratively efficient to apply it to the pool properties when all servers are expected to behave identically. Applying it to Nodes (Option C) would only verify that the IP address is up (typically via ICMP), which does not guarantee that the database service on port 1521 is actually responding. Finally, Profiles (Option A) are used to define how traffic is handled once it is accepted by a Virtual Server, not for the proactive health checking of backend resources.
Therefore, the pool properties page is the standard location for configuring service-specific availability requirements.

NEW QUESTION: 19
Refer to the exhibit.



A BIG-IP Administrator creates a new Virtual Server to load balance SSH traffic. Users are unable to log on to the servers.

What should the BIG-IP Administrator do to resolve the issue? (Choose one answer)

A. Set Protocol to UDP

B. Set Source Address to 10.1.1.2

C. Set Destination Address/Mask to 0.0.0.0/0

D. Set HTTP Profile to None

Answer: D (LEAVE A REPLY)

SSH is a Layer 4 TCP-based protocol that operates on TCP port 22 and does not use HTTP in any capacity. In the exhibit, the Virtual Server is configured with an HTTP Profile applied, which is inappropriate for SSH traffic and causes connection failures.

According to the BIG-IP Administration: Data Plane Configuration documentation:

An HTTP profile must only be applied to Virtual Servers handling HTTP or HTTPS traffic.

When an HTTP profile is attached, BIG-IP expects HTTP headers and attempts to parse application-layer data. Non-HTTP protocols such as SSH, FTP (control), SMTP, and other raw TCP services will fail if an HTTP profile is enabled.

Why the other options are incorrect:

A). Set Protocol to UDPSSH uses TCP, not UDP. Changing the protocol would break SSH entirely.

B). Set Source Address to 10.1.1.2The source address setting controls client access restrictions and is unrelated to protocol parsing issues.
C). Set Destination Address/Mask to 0.0.0.0/0The destination address is already valid for a specific SSH service and does not impact protocol handling.
Correct Resolution:
The BIG-IP Administrator should remove the HTTP Profile (set it to None) so the Virtual Server functions as a pure Layer 4 TCP service, allowing SSH connections to pass through successfully.

NEW QUESTION: 20
What is the status of a pool member when manual resume is enabled and a health check first fails and then passes?
A. Offline (Disabled)
B. Offline (Enabled)
C. Available (Disabled)
D. Available (Enabled)
Answer: A (LEAVE A REPLY)
With manual resume enabled, BIG-IP does not automatically return a pool member to service after recovery. The member remains offline until manually re-enabled.

NEW QUESTION: 21
In a pool there are 2 pool members (older servers) that can handle fewer connections than the other 3 newer servers.
Which load balancing method would allow more traffic to be directed to the newer servers? (Choose one answer)
A. Global Availability
B. Weighted Least Connections (member)
C. Round Robin
D. Least Connections (member)
Answer: B (LEAVE A REPLY)
This scenario requires unequal load distribution based on server capacity. The newer servers must receive more connections than the older ones, while still dynamically accounting for active connection counts.
According to BIG-IP Administration: Data Plane Configuration documentation:
Weighted Least Connections (member) combines:
Connection awareness (least connections)
Administrator-defined weights (ratios) to reflect server capacity
Pool members with higher weights receive proportionally more new connections than members with lower weights, even when using the same load balancing algorithm.
Why B is correct:
Allows assigning higher weights to newer servers and lower weights to older servers Ensures smarter traffic distribution based on both capacity and real-time load Why the other options are incorrect:
A). Global AvailabilityUsed for disaster recovery and site failover, not intra-pool load distribution.

C). Round RobinDistributes connections evenly without considering server capacity.
D). Least Connections (member)Balances only by current connection count and does not account for differences in server performance or capacity.
Correct Resolution:
Use Weighted Least Connections (member) and assign higher weights to newer servers so they receive more traffic while protecting older servers from overload.

NEW QUESTION: 22
In a pool there are 2 pool members (older servers) that can handle fewer connections than the other 3 newer servers.
Which load balancing method would allow more traffic to be directed to the newer servers? (Choose one answer)
A. Global Availability
B. Weighted Least Connections (member)
C. Round Robin
D. Least Connections (member)
Answer: B (LEAVE A REPLY)
This scenario requires unequal load distribution based on server capacity. The newer servers must receive more connections than the older ones, while still dynamically accounting for active connection counts.
According to BIG-IP Administration: Data Plane Configuration documentation:
Weighted Least Connections (member) combines:
Connection awareness (least connections)
Administrator-defined weights (ratios) to reflect server capacity
Pool members with higher weights receive proportionally more new connections than members with lower weights, even when using the same load balancing algorithm.
Why B is correct:
Allows assigning higher weights to newer servers and lower weights to older servers Ensures smarter traffic distribution based on both capacity and real-time load Why the other options are incorrect:
A . Global Availability
Used for disaster recovery and site failover, not intra-pool load distribution.
C . Round Robin
Distributes connections evenly without considering server capacity.
D . Least Connections (member)
Balances only by current connection count and does not account for differences in server performance or capacity.
Correct Resolution:
Use Weighted Least Connections (member) and assign higher weights to newer servers so they receive more traffic while protecting older servers from overload.

NEW QUESTION: 23
All pool members are online and all other settings are default.

What might alter the load balancing behavior?
A. Adding a OneConnect profile
B. Enabling SNAT Automap
C. Enabling an HTTP fallback host
D. Adding a persistence profile
Answer: (SHOW ANSWER)
Persistence overrides load balancing decisions by maintaining client-to-server affinity.

NEW QUESTION: 24
The BIG-IP Administrator has to provide encrypted communication between the users and the virtual server they access. Multiple hostnames are configured in DNS with the same IP address. Which profile type and setting in the profile should be used? (Choose one answer)
A. Client SSL, Client Name
B. Server SSL, Server Name
C. Client SSL, Server Name
D. Server SSL, Client Name
Answer: C (LEAVE A REPLY)
To provide encrypted communication between users and a virtual server, the BIG-IP system acts as a transparent SSL/TLS proxy. The administrative configuration required for this is a Client SSL profile.
When multiple hostnames (FQDNs) are associated with a single IP address, the system must determine which SSL certificate to present to the client during the initial TLS handshake. This is achieved using an extension of the TLS protocol called Server Name Indication (SNI).
The configuration logic is as follows:
* Profile Type: The Client SSL profile is responsible for terminating the SSL connection from the client (the user) to the BIG-IP system.
* The Setting: Within the Client SSL profile (under the "Advanced" view), there is a field specifically called Server Name. By entering the specific hostname (e.g., www.example.com) in this field, the BIG- IP system can match the hostname requested by the client in the ClientHello message to the correct profile.
* Implementation: The administrator typically creates multiple Client SSL profiles-one for each hostname-and assigns them all to the same virtual server. One of these profiles must be designated as the Default SSL Profile for SNI to handle requests where the client does not provide a hostname or provides one that does not match any specific profile.
By using the Client SSL profile and the Server Name setting, the BIG-IP system ensures that each user receives the correct certificate for the specific site they are trying to reach, even though all sites share a single virtual server IP.

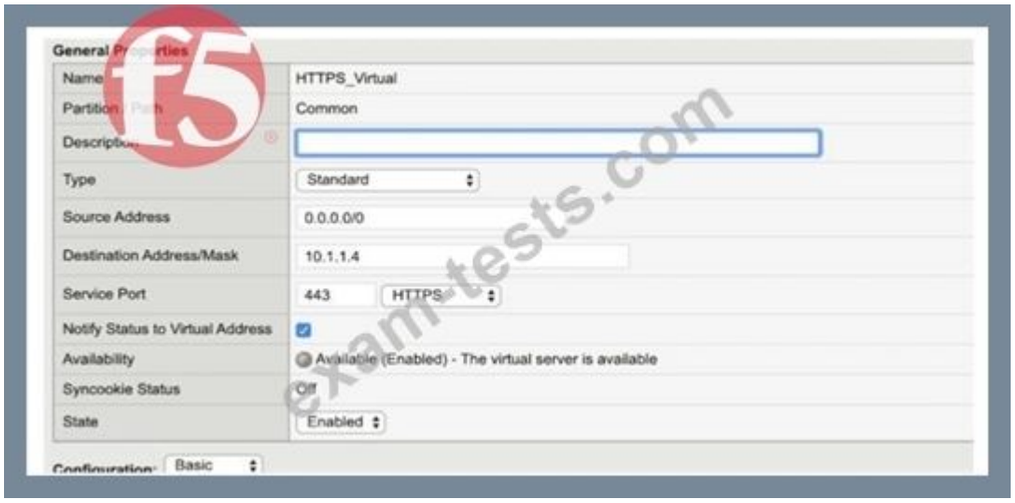
NEW QUESTION: 25
Refer to the exhibit.
A BIG-IP Administrator needs to configure health monitors for a newly configured server pool named Pool_B. Which health monitor settings will ensure that all pool members will be accurately marked as available or unavailable?

A. HTTPS, HTTP, FTP, and SSH, with the Availability Requirement of at least one monitor
B. HTTPS and HTTP with the Availability Requirement of at least one health monitor
C. HTTPS, HTTP, FTP, and ICMP, with the Availability Requirement of all health monitors
D. HTTPS, HTTP, FTP, and SSH with the Availability Requirement of all health monitors
Answer: (SHOW ANSWER)
The pool member list in the exhibit shows servers listening on four distinct service ports: 80 (HTTP), 21 (FTP), 443 (HTTPS), and 22 (SSH). To accurately monitor a pool where members provide different services, the administrator must apply monitors that correspond to each of those specific services.
The key to this configuration is the Availability Requirement (also known as " Monitor Rule "). If the administrator sets the requirement to " All health monitors, " every single monitor (HTTP, HTTPS, FTP, and SSH) must pass for a member to be marked " Up. " This would cause an immediate failure: for example, a server listening only on port 80 would fail the HTTPS, FTP, and SSH checks, resulting in the member being marked " Down " even if the HTTP service is healthy.
To ensure accuracy, the administrator should apply all four relevant monitors and set the Availability Requirement to at least one monitor. With this setting, the BIG-IP marks a member as " Up " if any of the assigned monitors return a successful response. Therefore, the member at 10.200.50.210:80 will stay " Up " as long as the HTTP monitor passes, even though it fails the FTP, HTTPS, and SSH monitors. This configuration allows a single pool to contain diverse service types while ensuring that the specific port defined for each member is verified correctly. Option C is incorrect because ICMP only checks if the IP is alive, not if the specific service port is functioning.

NEW QUESTION: 26
In the BIG-IP Configuration Utility, a user requests a single screen view to determine the status of all Virtual Servers and associated pool members, as well as any iRules in use. Where should the BIG-IP Administrator instruct the user to find this view?
A. Local Traffic > Monitors
B. Statistics
C. Local Traffic > Virtual Servers
D. Local Traffic > Network Map
Answer: (SHOW ANSWER)
The Local Traffic > Network Map provides a single consolidated view within the BIG-IP Configuration Utility that displays the full relationship hierarchy of all application delivery objects simultaneously. This includes Virtual Servers, their associated Pools, individual Pool Members, applied iRules, profiles, and the health status of each object - all rendered in a single topological screen.
This view is specifically designed for administrators and users who need rapid situational awareness across the data plane without navigating multiple menu screens. Each object is colour-coded to reflect its current availability state (green = available, red = unavailable, yellow = degraded).
The alternative options do not satisfy all requirements:
* Local Traffic > Monitors only displays health monitor configurations, not virtual server or pool member relationships.
* Statistics provides performance counters and throughput data, not object relationship mapping.

* Local Traffic > Virtual Servers lists virtual server configurations individually but does not provide a unified topology view inclusive of pool members and iRules in a single screen.
The Network Map is the authoritative single-pane-of-glass tool for application object visibility on BIG-IP.
Reference: BIG-IP Administration - Data Plane Configuration, Module: Monitoring and Managing Application Traffic - Network Map.

NEW QUESTION: 27
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 (Client)
B. Configure Protocol Profile (Server)
C. Configure SSL Profile (Server)
D. Configure Service Port to HTTP
Answer: (SHOW ANSWER)
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: 28
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? (Choose one answer)
A. Global Availability
B. Weighted Least Connections (member)
C. Round Robin
D. Least Connections (member)
Answer: B (LEAVE A REPLY)
When a pool contains servers with heterogeneous hardware capabilities (differing CPU, RAM, or connection limits), a static load balancing method like Round Robin is ineffective because it distributes requests equally, regardless of the server's capacity. To optimize traffic distribution for newer, more powerful servers, a dynamic or weighted method is required.
Weighted Least Connections (member): This is the ideal method for this scenario. It combines two factors:
Least Connections: It first checks the current active connection count to ensure traffic goes to the least busy server.
Weight (Ratio): It allows the administrator to assign a "Ratio" value to each pool member. Newer servers can be assigned a higher ratio (e.g., 3) while older servers are assigned a lower ratio (e.g., 1). The BIG-IP system uses these weights to disproportionately favor the newer servers even when connection counts are similar.
Why other options are incorrect:
Global Availability: This is primarily a GSLB (Global Server Load Balancing) or specific LTM priority group concept where traffic is sent to the first available member in a list until it fails, then moves to the next. It does not load balance based on capacity.
Round Robin: This passes each new connection request to the next server in line, treating the old and new servers exactly the same.
Least Connections (member): While this sends traffic to the server with the fewest active connections, it assumes all servers are equal. If an old server and a new server both have 10 connections, they are treated as equally capable of taking the 11th, which is not true in this scenario.

NEW QUESTION: 29
A BIG-IP Administrator creates a new VLAN on BIG-IP Cluster Member A and attaches an interface to it. Although the Auto Config Sync is in place, the new VLAN does NOT show up on Cluster Member B. What should the BIG-IP Administrator do to ensure the new VLAN is configured on each Cluster Member?
A. Reset the Device Trust of the BIG-IP Cluster on either Cluster Member
B. Configure a Default Route for the new VLAN on Cluster Member A
C. Configure the new VLAN manually on Cluster Member B
D. Enable the Interface that is attached to the new VLAN on Cluster Member A
Answer: C (LEAVE A REPLY)

NEW QUESTION: 30

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?

A. Auto Last Hop

B. Source Check

C. Fail-safe

D. Customer Tag

Answer: C (LEAVE A REPLY)

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.

NEW QUESTION: 31

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? (Choose one answer)

A. Global Availability

B. Weighted Least Connections (member)

C. Round Robin

D. Least Connections (member)

Answer: B (LEAVE A REPLY)

When a pool contains servers with heterogeneous hardware capabilities (differing CPU, RAM, or connection limits), a static load balancing method like Round Robin is ineffective because it distributes requests equally, regardless of the server's capacity. To optimize traffic distribution for newer, more powerful servers, a dynamic or weighted method is required.

* Weighted Least Connections (member): This is the ideal method for this scenario. It combines two factors:

* Least Connections: It first checks the current active connection count to ensure traffic goes to the least busy server.

* Weight (Ratio): It allows the administrator to assign a "Ratio" value to each pool member. Newer servers can be assigned a higher ratio (e.g., 3) while older servers are assigned a lower ratio (e.g.,

1). The BIG-IP system uses these weights to disproportionately favor the newer servers even when connection counts are similar.

* Why other options are incorrect:

* Global Availability: This is primarily a GSLB (Global Server Load Balancing) or specific LTM priority group concept where traffic is sent to the first available member in a list until it fails, then moves to the next. It does not load balance based on capacity.

* Round Robin: This passes each new connection request to the next server in line, treating the old and new servers exactly the same.

* Least Connections (member): While this sends traffic to the server with the fewest active connections, it assumes all servers are equal. If an old server and a new server both have 10 connections, they are treated as equally capable of taking the 11th, which is not true in this scenario.

Valid F5CAB3 Dumps shared by BraindumpsPass.com for Helping Passing F5CAB3 Exam! BraindumpsPass.com now offer the **newest F5CAB3 exam dumps**, the BraindumpsPass.com F5CAB3 exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com F5CAB3 dumps with Test Engine here: https://www.braindumpspass.com/F5/F5CAB3-practice-exam-dumps.html (**84** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 32

Which type of Virtual Server requires the use of a FastL4 profile?

A. Performance (Layer 4)

B. Stateless
C. Performance (HTTP)
D. Standard
Answer: (SHOW ANSWER)
Performance (Layer 4) virtual servers rely on FastL4 profiles for high-speed Layer 4 traffic handling.

NEW QUESTION: 33
An organization is reporting slow performance accessing their Intranet website. All employees use a single Proxy Server with a public IP.
What should the BIG-IP Administrator do to fix this issue?
A. Change Source Address to proxy IP
B. Change Load Balancing Method to Least Connections
C. Change Fallback Persistence Profile to source_addr
D. Change Default Persistence Profile to cookie
Answer: (SHOW ANSWER)
When multiple users share one source IP, source-address persistence fails. Cookie persistence uniquely identifies users at Layer 7 and ensures correct session handling.

NEW QUESTION: 34
An organization reports slow performance accessing an Intranet website. All employees use a single proxy IP.
What should the BIG-IP Administrator do?
A. Change Source Address to proxy IP
B. Change Load Balancing to Least Connections
C. Change Fallback Persistence to source_addr
D. Change Default Persistence to cookie
Answer: (SHOW ANSWER)
When many users share one source IP, source-address persistence fails. Cookie persistence uniquely identifies clients at Layer 7.

NEW QUESTION: 35
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: B,D (LEAVE A REPLY)
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.

NEW QUESTION: 36
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 (LEAVE A REPLY)
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.
While other options affect traffic handling, they do not "alter" the fundamental load balancing decision in the same way. A OneConnect profile (Option A) optimizes connection management by pooling idle server-side connections; while it changes how connections are reused, the initial load balancing decision still follows the configured method. A fallback host (Option C) is only utilized when the primary pool is unavailable, and since the question states all pool members are online, it remains inactive. SNAT Automap (Option D) changes the source IP address of the packet as it exits the BIG-IP toward the server to ensure return traffic passes back through the ADC, but it does not dictate which server is chosen for the request. Therefore, the persistence profile is the primary configuration element that overrides the load balancing algorithm to maintain a client-to-server relationship.

NEW QUESTION: 37
How can a BIG-IP Administrator identify a configuration object that has been configured within an iApp?
A. In both the GUI and the CLI the path will repeat the name of the object with .app appended to it, and the object type will be appended to the name. Example: ltm virtual /Common/testhttp.app/testhttp_vs

B. Using the CLI run the appropriate tmsh command to view the path and see the .app in it and the object type will be appended to the name. Example: ltm virtual /Common/testhttp.app/testhttp.vs
C. In both the GUI and the CLI the path will repeat the name of the object with .app appended to it and then name of the object. Example: ltm virtual /Common/testhttp.app/testhttp
Answer: (SHOW ANSWER)
In F5 BIG-IP systems, an iApp is a template-driven framework used to deploy and manage application- specific configurations. When an iApp is deployed, the system creates a specialized folder or " Application Service " container to hold all the resulting Local Traffic Manager (LTM) objects, such as virtual servers, pools, and profiles. This container is visually and programmatically identified by the .app extension in its path name.
The naming convention for these objects follows a specific hierarchical structure: / < partition > / < iapp_name > .app/ < object_name > . For example, if an administrator creates an iApp named " testhttp " in the " /Common " partition, the virtual server created by that iApp would be identified in the configuration (TMSH) and the GUI as /Common/testhttp.app/testhttp. This structure allows the BIG-IP system to distinguish between manually created objects and those managed by the iApp ' s automation logic.
This identification is crucial because iApp-managed objects typically have Strict Updates enabled, which prevents administrators from making direct manual changes to the object settings outside of the iApp ' s reconfigure interface. Attempting to modify such an object directly will result in an error message stating that the application service must be updated using an application management interface. By recognizing the .app path and the repeated name within the structure, an administrator can immediately identify that the object belongs to an iApp and should be managed through the iApp ' s specific management screen rather than standard LTM configuration menus.

NEW QUESTION: 38
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 (LEAVE A REPLY)
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.
While other tools have their place, they are inappropriate for this specific layer 3/4 connectivity issue. A qkview is a comprehensive diagnostic file used primarily for F5 Support to analyze the entire system's state but is

overkill for initial connectivity troubleshooting. An ssldump is used for inspecting SSL/TLS handshakes and encrypted payloads, which is irrelevant for a non-encrypted ICMP monitor. A curl command is a tool for testing HTTP/HTTPS application-level responses; it cannot be used to troubleshoot ICMP (ping) connectivity directly. By using tcpdump -ni <vlan_name> host <node_ip>, the administrator can see the ICMP "type 8" (request) and "type 0" (reply) packets, immediately identifying if the monitor failure is due to a "Destination Unreachable" message or a simple lack of response, thereby pinpointing the root cause in the data plane.

NEW QUESTION: 39
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. Ratio (member)
C. Round Robin
D. Weighted Least Connections (node)
Answer: (SHOW ANSWER)
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: 40
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 (LEAVE A REPLY)
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.

NEW QUESTION: 41
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 the GUI at the Resources tab for the virtual server.
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 TMSH with the list /ltm rule < irule > command.
Answer: (SHOW ANSWER)

NEW QUESTION: 42
A BIG-IP Administrator is creating a new Trunk on the BIG-IP device. What objects should be added to the new Trunk being created?
A. IP addresses
B. Network routes
C. VLANs
D. Interfaces
Answer: D (LEAVE A REPLY)

In F5 BIG-IP architecture, a Trunk is a logical aggregation of physical interfaces designed to increase bandwidth and provide link redundancy through 802.3ad Link Aggregation Control Protocol (LACP).
When creating a Trunk object within the BIG-IP configuration, the administrator must assign physical Interfaces to it - these are the underlying hardware ports (e.g., 1.1, 1.2, 1.3) that will be bundled together to form the aggregated logical link.
It is critical to understand the layered object hierarchy on the BIG-IP platform:
* Interfaces # assigned to Trunks
* Trunks (or individual Interfaces) # assigned to VLANs (as tagged or untagged members)
* VLANs # assigned Self IP addresses
* Self IPs / Routes # configured at the network layer above
VLANs are not added to a Trunk; rather, a Trunk is added as a member of a VLAN. IP addresses and network routes are configured independently at a higher abstraction layer and have no direct association with Trunk creation.
This layered design ensures proper separation of concerns across the data plane. Misunderstanding this hierarchy is a common configuration error that can result in traffic forwarding failures.
Reference: BIG-IP Administration - Data Plane Configuration, Module: Network Objects - Interfaces, Trunks, and VLANs.

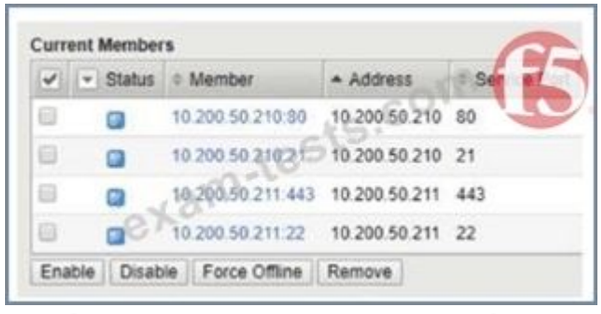
NEW QUESTION: 43
A BIG-IP Administrator is setting up a new BIG-IP device. The network administrator reports that the interface has an incompatible media speed. The BIG-IP Administrator needs to change this setting manually.
From which location should the BIG-IP Administrator perform this task?
A. On the Front Console
B. In the TMOS Shell Command line
C. In the Configuration Utility, Network > Interface
D. In the Configuration Utility, System > Configuration
Answer: C (LEAVE A REPLY)
Standard BIG-IP administration dictates that hardware-level physical attributes are managed within the Network section of the configuration. When a network switch and a BIG-IP fail to successfully negotiate speed and duplex settings (Auto-Negotiation), it can result in CRC errors, late collisions, or a total lack of link. To resolve this manually, the administrator must navigate to the Configuration Utility (GUI) and go to Network > Interfaces. Within the Interfaces list, the administrator can select the specific physical port (e.g., 1.1 or 1.2) and modify its properties. By default, the media speed is set to "Auto," but the drop-down menu allows for manual selection of specific speeds (e.g., 100Mb/s, 1Gb/s, 10Gb/s) and duplex settings (Full or Half). While these changes can also be made via the TMOS Shell (TMSH) (Option B) using the modify net interface command, the question asks for the standard location, which in most administrative contexts refers to the primary GUI path. System > Configuration (Option D) is used for global device settings like NTP, DNS, and licensing, not for interface-specific physical layer parameters. The Front Console (Option A), referring to the LCD panel on physical appliances, is primarily used for initial management IP setup and viewing system alerts, but does not provide the granular interface configuration required for media speed adjustments.

NEW QUESTION: 44
In a pool there are 2 pool members (older servers) that can handle fewer connections than the other 3 newer servers.
Which load balancing method would allow more traffic to be directed to the newer servers? (Choose one answer)
A. Global Availability
B. Weighted Least Connections (member)
C. Round Robin
D. Least Connections (member)
Answer: B (LEAVE A REPLY)
This scenario requires unequal load distribution based on server capacity. The newer servers must receive more connections than the older ones, while still dynamically accounting for active connection counts.
According to BIG-IP Administration: Data Plane Configuration documentation:
* Weighted Least Connections (member) combines:
* Connection awareness (least connections)
* Administrator-defined weights (ratios) to reflect server capacity
* Pool members with higher weights receive proportionally more new connections than members with lower weights, even when using the same load balancing algorithm.
Why B is correct:
* Allows assigning higher weights to newer servers and lower weights to older servers
* Ensures smarter traffic distribution based on both capacity and real-time load Why the other options are incorrect:
* A. Global AvailabilityUsed for disaster recovery and site failover, not intra-pool load distribution.
* C. Round RobinDistributes connections evenly without considering server capacity.
* D. Least Connections (member)Balances only by current connection count and does not account for differences in server performance or capacity.
Correct Resolution:
Use Weighted Least Connections (member) and assign higher weights to newer servers so they receive more traffic while protecting older servers from overload.

NEW QUESTION: 45
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? (Choose two answers)
A. Via TMSH with the list /ltm rule <irule> command.
B. Via the GUI at the Resources tab for the virtual server.
C. Via TMSH with the list /ltm virtual <virtual_server> command.
D. Via the GUI at the iRule tab for the virtual server.
Answer: B,C (LEAVE A REPLY)

In BIG-IP systems, iRules influence traffic only when they are attached to a Virtual Server. If application traffic is being sent to nodes or pool members that are not defined in the pool, this typically indicates that an iRule is overriding the default load-balancing behavior by explicitly selecting a pool or node.
According to BIG-IP Administration: Data Plane Configuration and official F5 guidance:
* iRules are associated with Virtual Servers, not directly with pools or nodes.
* To determine whether an iRule is actively affecting traffic, the administrator must inspect the Virtual Server configuration.
Explanation of the correct answers:
* B. Via the GUI at the Resources tab for the virtual serverThe Resources tab in the Configuration Utility displays all traffic-handling objects applied to the Virtual Server, including assigned iRules. This is the primary GUI location to verify whether an iRule is influencing data plane traffic.
* C. Via TMSH with the list /ltm virtual <virtual_server> commandThis TMSH command displays the full Virtual Server configuration, including any iRules listed under the rules section. It is the authoritative CLI method to confirm iRule usage.
Why the other options are incorrect:
* A. Via TMSH with the list /ltm rule <irule> commandThis command only shows the contents of an iRule and does not indicate whether the iRule is attached to or used by any Virtual Server.
* D. Via the GUI at the iRule tab for the virtual serverBIG-IP does not provide a dedicated "iRule" tab on Virtual Servers. iRules are viewed and managed under the Resources tab.
Correct Conclusion:
To verify whether an iRule is responsible for unexpected node selection, the BIG-IP Administrator must examine the Virtual Server configuration, either through the Resources tab in the GUI or by using TMSH to list the Virtual Server configuration.

NEW QUESTION: 46
Refer to the exhibit.



A BIG-IP Administrator needs to configure health monitors for a newly configured server pool named Pool_B. Which health monitor settings will ensure that all pool members will be accurately marked as available or unavailable?
A. HTTPS, HTTP, FTP, and SSH, with the Availability Requirement of at least one monitor
B. HTTPS and HTTP with the Availability Requirement of at least one health monitor
C. HTTPS, HTTP, FTP, and ICMP, with the Availability Requirement of all health monitors
D. HTTPS, HTTP, FTP, and SSH with the Availability Requirement of all health monitors
Answer: A (LEAVE A REPLY)

The pool member list in the exhibit shows servers listening on four distinct service ports: 80 (HTTP), 21 (FTP), 443 (HTTPS), and 22 (SSH). To accurately monitor a pool where members provide different services, the administrator must apply monitors that correspond to each of those specific services.
The key to this configuration is the Availability Requirement (also known as "Monitor Rule"). If the administrator sets the requirement to "All health monitors," every single monitor (HTTP, HTTPS, FTP, and SSH) must pass for a member to be marked "Up." This would cause an immediate failure: for example, a server listening only on port 80 would fail the HTTPS, FTP, and SSH checks, resulting in the member being marked "Down" even if the HTTP service is healthy.
To ensure accuracy, the administrator should apply all four relevant monitors and set the Availability Requirement to at least one monitor. With this setting, the BIG-IP marks a member as "Up" if any of the assigned monitors return a successful response. Therefore, the member at 10.200.50.210:80 will stay "Up" as long as the HTTP monitor passes, even though it fails the FTP, HTTPS, and SSH monitors. This configuration allows a single pool to contain diverse service types while ensuring that the specific port defined for each member is verified correctly. Option C is incorrect because ICMP only checks if the IP is alive, not if the specific service port is functioning.

Valid F5CAB3 Dumps shared by BraindumpsPass.com for Helping Passing F5CAB3 Exam! BraindumpsPass.com now offer the **newest F5CAB3 exam dumps**, the BraindumpsPass.com F5CAB3 exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com F5CAB3 dumps with Test Engine here: https://www.braindumpspass.com/F5/F5CAB3-practice-exam-dumps.html (**84** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)

NEW QUESTION: 47
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 (LEAVE A REPLY)
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.

NEW QUESTION: 48
How can a BIG-IP Administrator identify a configuration object that has been configured within an iApp?
A. In both the GUI and the CLI the path will repeat the name of the object with .app appended to it, and the object type will be appended to the name. Example: ltm virtual /Common/testhttp.app/testhttp_vs
B. Using the CLI run the appropriate tmsh command to view the path and see the .app in it and the object type will be appended to the name. Example: ltm virtual /Common/testhttp.app/testhttp.vs
C. In both the GUI and the CLI the path will repeat the name of the object with .app appended to it and then name of the object. Example: ltm virtual /Common/testhttp.app/testhttp
Answer: (SHOW ANSWER)
In F5 BIG-IP systems, an iApp is a template-driven framework used to deploy and manage application- specific configurations. When an iApp is deployed, the system creates a specialized folder or "Application Service" container to hold all the resulting Local Traffic Manager (LTM) objects, such as virtual servers, pools, and profiles. This container is visually and programmatically identified by the .app extension in its path name.
The naming convention for these objects follows a specific hierarchical structure: /<partition>/<iapp_name>.app/<object_name>. For example, if an administrator creates an iApp named "testhttp" in the "/Common" partition, the virtual server created by that iApp would be identified in the configuration (TMSH) and the GUI as /Common/testhttp.app/testhttp. This structure allows the BIG-IP system to distinguish between manually created objects and those managed by the iApp's automation logic.
This identification is crucial because iApp-managed objects typically have Strict Updates enabled, which prevents administrators from making direct manual changes to the object settings outside of the iApp's reconfigure interface. Attempting to modify such an object directly will result in an error message stating that the application service must be updated using an application management interface. By recognizing the .app path and the repeated name within the structure, an administrator can immediately identify that the object belongs to an iApp and should be managed through the iApp's specific management screen rather than standard LTM configuration menus.

Valid F5CAB3 Dumps shared by BraindumpsPass.com for Helping Passing F5CAB3 Exam! BraindumpsPass.com now offer the **newest F5CAB3 exam dumps**, the BraindumpsPass.com F5CAB3 exam **questions have been updated** and **answers have been corrected** get the **newest** BraindumpsPass.com F5CAB3 dumps with Test Engine here: https://www.braindumpspass.com/F5/F5CAB3-practice-exam-dumps.html (**84** Q&As Dumps, **40%OFF** **Special Discount: Exam-Tests**)