F5.F5CAB2.v2026-03-13.q19

Exam Code:	F5CAB2
Exam Name:	BIG-IP Administration Data Plane Concepts (F5CAB2)
Certification Provider:	F5
Free Question Number:	19
Version:	v2026-03-13
# of views:	107
# of Questions views:	190
https://www.freecram.net/torrent/F5.F5CAB2.v2026-03-13.q19.html	

NEW QUESTION: 1
The BIG-IP Administrator wants to provide quick failover between the F5 LTM devices that are configured as an HA pair with a single-selfip using the MAC Masquerade feature for this quick failover and runs this command: tmsh modify /cm traffic-group traffic-group-1 mac 02:12:34:56:00:00 However, the Network Operations team has identified an issue with the use of the same MAC address being used within different VLANs. As a result, the administrator decides to implement the Per-VLAN Mac Masquerade in order to have a unique MAC address on each VLAN: tmsh modify /sys db tm.macmasqaddr_per_vlan value true. What would be the resulting MAC address on a tagged VLAN of 1501? (Choose one answer)
A. 02:12:34:56:01:15
B. 02:12:34:56:dd:05
C. 02:12:34:56:05:dd
D. 02:12:34:56:15:01
Answer: C (LEAVE A REPLY)
According to F5 BIG-IP documentation regarding High Availability and MAC Masquerade behavior, the system allows for more granular control over Layer 2 addresses during failover events.
* Standard MAC Masquerade: By default, when a traffic group is assigned a MAC masquerade address (like 02:12:34:56:00:00), the BIG-IP system uses that exact MAC address for all traffic associated with that traffic group across all VLANs. This ensures that upstream switches do not need to relearn ARP entries for the Virtual IP, but it can cause issues in environments where multiple VLANs share the same physical infrastructure or monitoring tools that flag identical MACs across segments.
* Per-VLAN MAC Masquerade: When the system database variable tm.macmasqaddr_per_vlan is set to true, the BIG-IP system calculates a unique MAC address for each VLAN. It does this by taking the base MAC masquerade address configured in the traffic group and adding the VLAN ID (tag) to it.

* Calculation Logic:
* Base MAC: 02:12:34:56:00:00
* VLAN ID: 1501
* To find the suffix, the VLAN ID is converted from decimal to hexadecimal:
* 1501 in decimal = 05DD in hex.
* The system then applies this offset to the last two octets of the base MAC address.
* 00:00 + 05:DD = 05:DD.
* Result: The final MAC address for VLAN 1501 becomes 02:12:34:56:05:dd.
This ensures that every VLAN has a unique Layer 2 identity while still reaping the benefits of "gratuitous ARP-less" failover provided by MAC masquerading.

NEW QUESTION: 2
What type of virtual server should be used to block responses for one IP in a subnet with a virtual server?
(Choose one answer)
A. Drop
B. Block
C. Reject
D. Standard
Answer: (SHOW ANSWER)
In the BIG-IP system, when you need to prevent traffic from reaching a specific destination or being processed by the system, you utilize specific Virtual Server types that act as "denial" points.
* Reject Virtual Servers: When a packet matches a Reject virtual server, the BIG-IP system stops the packet from being processed and sends a reset (RST) in the case of TCP, or an ICMP unreachable message in the case of UDP. This is the preferred method for "blocking" specific IPs when you want the sender to receive immediate notification that the connection was refused.
* Drop Virtual Servers: A Drop virtual server simply discards the packet without sending any response back to the source. While effective for "stealthing" a network, it is often less desirable for standard administration unless specifically mitigating a DoS attack.
* Comparison with Standard: A Standard virtual server is used to process and load balance traffic to a pool of members; it does not inherently act as a "blocking" mechanism for a single IP within a subnet unless combined with complex iRules or Packet Filters.
* Context of the Question: To block responses (or connection attempts) for a specific IP while other traffic in the subnet might be handled by more permissive virtual servers, a more specific (higher precedence) Reject virtual server is the standard administrative approach.

NEW QUESTION: 3
To increase the available bandwidth of an existing trunk, the BIG-IP Administrator plans to add additional interfaces. Which command should the BIG-IP Administrator run from within the bash shell? (Choose one answer)
A. tmsh modify /net trunk trunk_A interfaces add {1.3 1.4}

B. tmsh modify /sys trunk trunk_A interfaces add {1.3 1.4}
C. tmsh create /net trunk trunk_A interfaces add {1.3 1.4}
D. tmsh create /sys trunk trunk_A interfaces add {1.3 1.4}
Answer: (SHOW ANSWER)
In BIG-IP, a trunk is a Layer 2 network object used to aggregate multiple physical interfaces into a single logical link. This aggregation provides increased bandwidth and link resiliency, commonly in conjunction with LACP.
Key concepts that apply here:
* Trunks are managed under the /net trunk tmsh hierarchy
* Physical interfaces are added or removed using the modify command
* The create command is used only when defining a brand-new trunk, not when updating an existing one Because the trunk already exists and the goal is to add interfaces, the correct operation is:
tmsh modify /net trunk trunk_A interfaces add {1.3 1.4}
This command:
* Modifies the existing trunk named trunk_A
* Adds interfaces 1.3 and 1.4 to the trunk
* Immediately increases available bandwidth and redundancy
Why the Other Options Are Incorrect
* B uses the /sys hierarchy, which is not used for trunks
* C attempts to create a trunk that already exists
* D uses an incorrect hierarchy and an incorrect operation

NEW QUESTION: 4
An ecommerce company is experiencing latency issues with online shops during Black Friday's peak season.
The BIG-IP Administrator detects an overall high CPU load on the BIG-IP device and wants to move the top utilized Virtual Servers to a dedicated BIG-IP device. Where should the BIG-IP Administrator determine the problematic Virtual Servers? (Choose one answer)
A. Local Traffic > Virtual Servers > Virtual Server List
B. System > Platform
C. Statistics > Module Statistics > Local Traffic > Virtual Servers
D. Local Traffic > Network Map
Answer: (SHOW ANSWER)
When troubleshooting performance and latency issues on BIG-IP, especially under peak load conditions, it is critical to identify which Virtual Servers are consuming the most resources. This is a core data plane analysis task.
BIG-IP provides multiple views of configuration and status, but only certain areas expose real-time and historical traffic statistics that correlate directly with CPU usage and throughput.
Why Option C Is Correct:
Statistics > Module Statistics > Local Traffic > Virtual Servers provides:

* Real-time and cumulative statistics per Virtual Server
* Metrics such as:
* Bits in / Bits out
* Packets in / Packets out
* Current connections
* Connection rate
* Total requests
* The ability to identify high-traffic or high-connection Virtual Servers, which are the most likely contributors to elevated CPU utilization These statistics allow the administrator to objectively determine which Virtual Servers are the top consumers of system resources and therefore good candidates for migration to a dedicated BIG-IP device.
Why the Other Options Are Incorrect:
* A. Local Traffic > Virtual Servers > Virtual Server List
* Primarily a configuration view
* Does not provide sufficient performance or utilization statistics to identify CPU-heavy Virtual Servers
* B. System > Platform
* Displays hardware-level information such as CPU cores, memory, disk, and platform type
* Does not break down utilization by Virtual Server
* D. Local Traffic > Network Map
* Provides a logical topology view of Virtual Servers, pools, and pool members
* Useful for understanding relationships, but not for identifying high-utilization Virtual Servers Key Data Plane Concept Reinforced:
To diagnose performance problems and plan traffic redistribution, BIG-IP administrators must rely on Module and object-level statistics, not configuration screens. The Virtual Server statistics view is the authoritative location for identifying traffic hotspots that directly impact CPU and latency during peak events such as Black Friday.

NEW QUESTION: 5
A BIG-IP Administrator needs to connect a BIG-IP system to two upstream switches to provide external network resilience. The network engineer instructs the administrator to configure interface binding with LACP. Which configuration should the administrator use? (Choose one answer)
A. A virtual server with an LACP profile and the switches' management IPs as pool members.
B. A virtual server with an LACP profile and the interfaces connected to the switches as pool members.
C. A Trunk listing the allowed VLAN IDs and MAC addresses configured on the switches.
D. A Trunk containing an interface connected to each switch.
Answer: (SHOW ANSWER)
In BIG-IP architecture, link aggregation and redundancy at Layer 2 are implemented using Trunks, not virtual servers or pools.

According to BIG-IP Administration Data Plane Concepts:
* Interfaces are the physical network ports on the BIG-IP device
* A Trunk is a logical grouping of multiple interfaces
* Trunks can be configured to use LACP (Link Aggregation Control Protocol) to:
* Provide link redundancy
* Increase aggregate bandwidth
* Allow automatic detection of link failures
* VLANs are then assigned to the trunk, not directly to individual interfaces, once aggregation is in place Correct Design for the Scenario:
To connect BIG-IP to two upstream switches with LACP:
* One physical interface from BIG-IP connects to Switch A
* Another physical interface from BIG-IP connects to Switch B
* Both interfaces are placed into the same trunk
* LACP is enabled on the trunk and on the switches
This configuration allows:
* Traffic to continue flowing if one interface or switch fails
* Proper LACP negotiation between BIG-IP and the upstream switches
* Clean separation of responsibilities (Layer 2 handled by trunking, Layer 4-7 by virtual servers)
Why Option D Is Correct:
* A Trunk containing an interface connected to each switch is exactly how BIG-IP implements LACP- based interface binding
* The trunk handles link state, load distribution, and failover at the data plane Why the Other Options Are Incorrect:
* A & B - Virtual servers operate at Layers 4-7 and have nothing to do with physical link aggregation or LACP
* C - VLAN IDs and MAC addresses are not configured inside a trunk definition; trunks aggregate interfaces, and VLANs are applied to trunks Key Data Plane Concept Reinforced:
On BIG-IP systems, LACP is always configured on a Trunk, which aggregates physical interfaces to provide Layer 2 resiliency and bandwidth aggregation. Virtual servers and pools are not involved in physical interface binding.

NEW QUESTION: 6
The diagram below shows the TCP connection setup for an application.

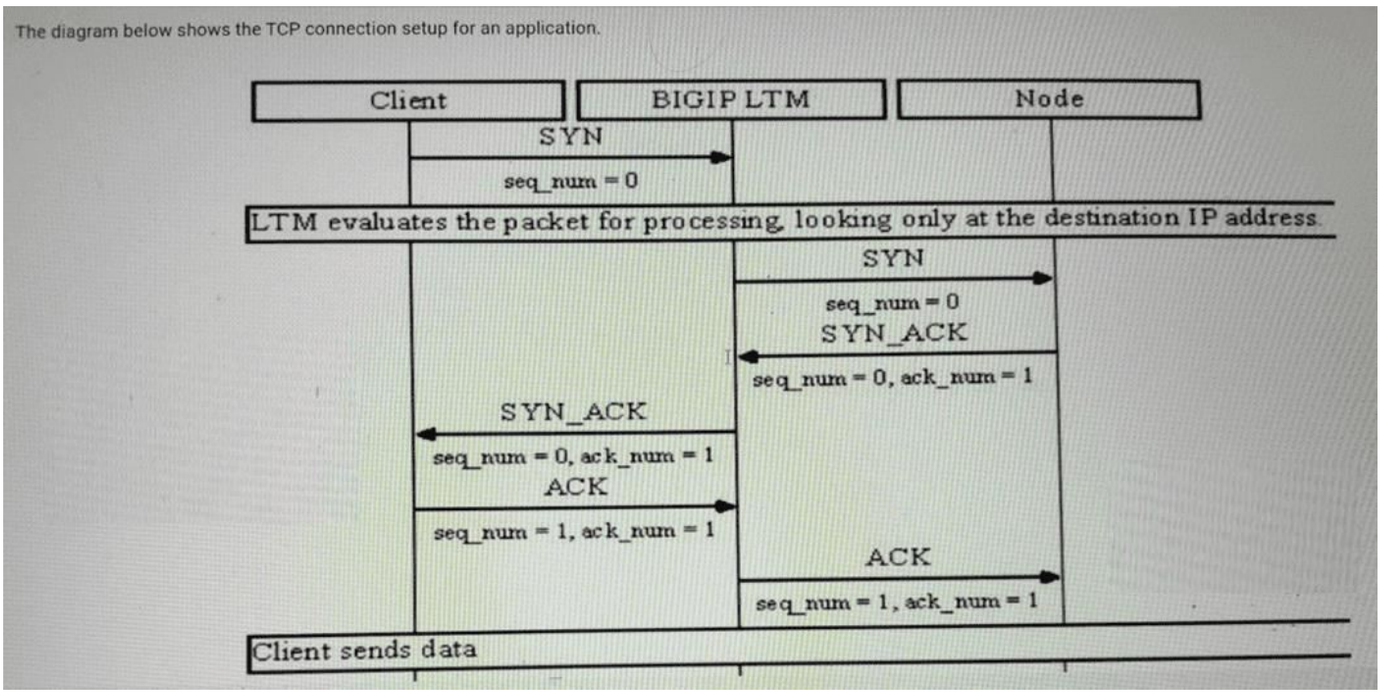

Which of the following virtual server types applies? (Choose one answer)

A. Standard virtual server

B. Forwarding IP virtual server

C. Stateless virtual server

Answer: B (LEAVE A REPLY)

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

* Initial Packet Processing: The diagram explicitly states that the LTM evaluates the packet looking only at the destination IP address. This is the fundamental characteristic of a Forwarding IP virtual server, which uses the system's routing table to make forwarding decisions instead of load balancing to a pool of members.
* Handshake Sequence: Unlike a Standard virtual server, which completes the three-way handshake with the client (SYN, SYN-ACK, ACK) before initiating a separate connection to the server, the Forwarding IP virtual server passes the client's original SYN packet directly to the destination node.
* Response Timing: The BIG-IP system waits for the SYN-ACK from the destination node before it sends a SYN-ACK back to the client. It essentially "passes through" the handshake signals while still maintaining a state entry in the connection table to track the flow.
* Packet-by-Packet Logic: While it tracks the state, it does not perform address translation (unless SNAT is specifically configured) or deep packet inspection like a full proxy would.

Why other options are incorrect:

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

* Stateless virtual server: A stateless virtual server does not track connections in the connection table.
The diagram shows the system meticulously passing sequence numbers (seq_num) and acknowledgment numbers (ack_num) between the two sides, which requires stateful tracking of the TCP flow.

NEW QUESTION: 7
Which three iRule events are likely to be seen in iRules designed to select a pool for load balancing? (Choose three.)
A. CLIENT_DATA
B. SERVER_DATA
C. HTTP_REQUEST
D. HTTP_RESPONSE
E. CLIENT_ACCEPTED
F. SERVER_SELECTED
G. SERVER_CONNECTED
Answer: A,C,E (LEAVE A REPLY)
12
In the BIG-IP system, pool selection must occur on the client-side of the connection, before the system attempts to connect to a pool 3member. The events listed 4are the primary entry points for making these decisions:
* CLIENT_ACCEPTED (E): This is a Layer 4 event triggered when the BIG-IP accepts a TCP connection. It is the earliest point where a pool can be assigned based on the client's source IP address or the destination port.
* CLIENT_DATA (A): This event is triggered when the system receives a "chunk" of data on the client- side. It is often used for non-HTTP protocols (like custom TCP protocols) to inspect the payload and select a pool based on its contents.
* HTTP_REQUEST (C): This is a Layer 7 event. It occurs once the BIG-IP has fully parsed the HTTP headers. This is the most common event for pool selection, allowing the administrator to route traffic based on the URI, Host header, or cookies.
Events like SERVER_SELECTED or SERVER_CONNECTED occur after the load balancing decision has already been made, and HTTP_RESPONSE or SERVER_DATA occur after the server has already started communicating back, making them too late for initial pool selection.

NEW QUESTION: 8
What is required for a virtual server to support clients whose traffic arrives on the internal VLAN and pool members whose traffic arrives on the external VLAN?
A. That support is never available.
B. The virtual server must be enabled for both VLANs.
C. The virtual server must be enabled on the internal VLAN.
D. The virtual server must be enabled on the external VLAN.

Answer: C (LEAVE A REPLY)
4647
Virtual Servers have a setting called VLAN and Tunnel Traffic which defines where the BIG-IP "listens" for new connections.4849
* Ingress Logic: A virtual server is an entry point. It must be enab50led on the VLAN where the Client resides. If a client is on the "51Internal" VLAN, the Virtual Server must be enabled there to receive the traffic.
* Egress Logic: The BIG-IP system uses the TMM Routing Table and Self-IPs to reach pool members.
It does not need the Virtual Server to be "enabled" on the destination VLAN (External) to send traffic there.
* Default Behavior: By default, Virtual Servers are enabled on "All VLANs." However, if restricted for security, the administrator must ensure the Virtual Server is active on the client-facing (ingress) VLAN.

NEW QUESTION: 9
A BIG-IP Administrator makes a configuration change to a Virtual Server on the Standby device of an HA pair. The HA pair is currently configured with Auto-Sync enabled. What effect will the change have on the HA pair configuration? (Choose one answer)
A. The change will be undone next time a configuration change is made on the Active device.
B. The change will be propagated next time a configuration change is made on the Active device.
C. The change will be undone when Auto-Sync propagates the config to the Standby device.
D. The change will take effect when Auto-Sync propagates the config to the HA pair.
Answer: (SHOW ANSWER)
In a BIG-IP high availability (HA) configuration, Auto-Sync is a device trust feature that automatically synchronizes configuration changes from the Active device to the Standby device within a Sync-Failover device group.
Key principles from BIG-IP Administration Data Plane Concepts:
* The Active device is always the authoritative source of configuration
* Configuration changes are intended to be made only on the Active device
* With Auto-Sync enabled, any time the Active device configuration changes, the system automatically pushes the configuration to all Standby members of the device group
* Configuration changes made directly on a Standby device are not preserved In this scenario:
* The administrator modifies a Virtual Server on the Standby device
* That change is local only and does not alter the device group's synchronized configuration
* When Auto-Sync next runs (triggered by a change on the Active device or an internal sync event), the Active device configuration overwrites the Standby configuration As a result, the configuration change made on the Standby device is undone.
Why the Other Options Are Incorrect:
* A - The change is not undone only when another change is made; it is undone during the next Auto- Sync operation

* B - Changes made on the Standby device are never propagated to the Active device
* D - Auto-Sync does not merge or promote Standby changes into the HA pair configuration Best Practice Reinforced:
Always perform configuration changes on the Active BIG-IP device when Auto-Sync is enabled to ensure consistent and predictable HA behavior.

NEW QUESTION: 10

A. Two members offline and one member online
B. Two members online and one member offline
C. All members offline
D. All members online
Answer: (SHOW ANSWER)
In BIG-IP LTM, health monitors are used to determine the availability of pool members and directly influence traffic flow decisions in the data plane.
Key characteristics of the default HTTP monitor according to BIG-IP Administration Data Plane Concepts:
* Sends an HTTP request (typically GET /)
* Expects an HTTP response code of 200 OK
* Any response other than 200 is treated as a monitor failure
* A failed monitor causes the pool member to be marked offline (down)
In this scenario:
* Two pool members return 404 Not Found
* A 404 response indicates that the requested object was not found
* This does not meet the success criteria of the default HTTP monitor
* These two members are therefore marked offline
* One pool member returns 200 OK
* This matches the expected response
* The member is marked online
Resulting Pool Member Availability:
* 2 members: Offline
* 1 member: Online
Why the Other Options Are Incorrect:
* B - 404 responses are not considered healthy by the default HTTP monitor
* C - At least one member responds with the expected 200 OK
* D - Members returning 404 responses fail the monitor and cannot be marked online Key Data Plane Concept Reinforced:
BIG-IP health monitors make binary availability decisions based strictly on configured success criteria. For HTTP monitors, response codes matter-404 is a failure, even if the service is technically reachable.

NEW QUESTION: 11

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

NEW QUESTION: 12
A virtual server is listening at 10.10.1.100:80 and has the following iRule associated with it:
when HTTP_REQUEST { if { [HTTP::header UserAgent] contains "MSIE" }
{ pool MSIE_pool }
else { pool Mozilla_pool }
If
a user connects to http://10.10.1.100/foo.html and their browser does not specify a UserAgent, which pool will receive the request?

A. None. The request will be dropped.
B. Mozilla_pool
C. MSIE_pool
D. Unknown. The pool cannot be determined from the information provided.
Answer: (SHOW ANSWER)

NEW QUESTION: 13
A BIG-IP system receives UDP traffic from a specific source. The administrator wants the traffic to be forwarded, not dropped or rejected. Which virtual server type should be used? (Choose one answer)
A. Drop
B. Reject
C. Block
D. Standard
Answer: (SHOW ANSWER)
BIG-IP virtual server types define how traffic is handled at the data plane when it matches a virtual server's destination address and service port.
According to BIG-IP Administration Data Plane Concepts:
* Standard virtual server
* The default and most commonly used type
* Accepts client connections and forwards traffic to pool members
* Supports both TCP and UDP traffic
* Allows full use of profiles (UDP, FastL4, persistence, etc.) and iRules
* Required when the goal is to process and pass traffic through BIG-IP
* Drop virtual server
* Silently discards matching traffic
* No response is sent to the client
* Reject virtual server
* Actively rejects traffic by sending an error response
* For UDP, BIG-IP may send an ICMP unreachable message
* Block virtual server
* Used to block traffic at the virtual server level
* Traffic is neither forwarded nor processed by pools
In this scenario:
* The administrator explicitly wants the UDP traffic to be forwarded
* Only a Standard virtual server forwards traffic to a pool or next-hop destination Why the Other Options Are Incorrect:
* A. Drop - Traffic is silently discarded
* B. Reject - Traffic is actively rejected
* C. Block - Traffic is blocked and not forwarded
Key Data Plane Concept Reinforced:

When traffic must be accepted and forwarded-regardless of whether it is TCP or UDP-the BIG-IP administrator must use a Standard virtual server, which is the only virtual server type designed for normal application traffic processing.

NEW QUESTION: 14
When upgrading a BIG-IP redundant pair, what happens when one system has been updated but the other has not?
A. Synching should not be performed.
B. The first system to be updated will assume the Active role.
C. This is not possible since both systems are updated simultaneously.
D. The older system will issue SNMP traps indicating a communication error with the partner.
Answer: (SHOW ANSWER)
The F5 BIG-IP upgrade process for HA pairs requires a specific "staggered" approach to maintain uptime.
* Version Mismatch: When one unit is upgraded to a newer version of TMOS (e.g., from 15.1 to 16.1), it enters a26 "Version Mismatch" 27state with its peer.
* Configuration Sync: Because the configuration schemas between different versions are often incompatible, ConfigSync should not be performed. Attempting to sync a newer configuration to an older system (or vice-versa) can cause configuration corruption or system instability.
* Failover Capability: Generally, a pair with a version mismatch can still fail over to ensure traffic continuity during the upgrade window, but administrative changes and syncs must be paused until both units are on the same version.

NEW QUESTION: 15
A BIG-IP Administrator has a cluster of devices.
What should the administrator do after creating a new Virtual Server on device 1? (Choose one answer)
A. Synchronize the settings of the group to device 1
B. Create a new cluster on device 1
C. Synchronize the settings of device 1 to the group
D. Create a new virtual server on device 2
Answer: (SHOW ANSWER)
In a BIG-IP device service cluster, configuration objects such as virtual servers, pools, profiles, and iRules are maintained through configuration synchronization (config-sync).
Key BIG-IP concepts involved:
* Device Service Cluster (DSC)A cluster is a group of BIG-IP devices that share configuration data. One device is typically used to make changes, which are then synchronized to the rest of the group.
* Config-Sync Direction Matters
* Changes are made on a local device
* Those changes must be pushed to the group

* The correct operation is "Sync Device to Group"
Why C is correct:
* The virtual server was created only on device 1
* Other devices in the cluster do not yet have this object
* To propagate the new virtual server to all cluster members, the administrator must synchronize device
1 to the group
Why the other options are incorrect:
* A. Synchronize the settings of the group to device 1This would overwrite device 1's configuration with the group's existing configuration and may remove the newly created virtual server.
* B. Create a new cluster on device 1The cluster already exists. Creating a new cluster is unnecessary and disruptive.
* D. Create a new virtual server on device 2This defeats the purpose of centralized configuration management and risks configuration drift.
Conclusion:
After creating a new virtual server on a BIG-IP device that is part of a cluster, the administrator must synchronize the configuration from that device to the group so all devices share the same ADC application objects.

NEW QUESTION: 16
When using the setup utility to configure a redundant pair, you are asked to provide a "Failover Peer IP".
Which address is this?
A. an address of the other system in its management network
B. an address of the other system in a redundant pair configuration
C. an address on the current system used to listen for failover messages from the partner BIG-IP
D. an address on the current system used to initiate mirroring and network failover heartbeat messages
Answer: (SHOW ANSWER)
When establishing a redundant pair, each device must know where to send its health heartbeats and sync data.
* The Peer IP: The Failover Peer IP is the IP address belonging to the other BIG-IP device in the HA pair. This is typically a34 Self-IP on a dedicated "HA" or "Internal" VLAN, or the Management IP.
* Purpose: It identifies the destination for the "Heartbeat" (the "Are you alive?" check).
* Setup Context: During the initial setup, you tell Device A to look for Device B at its "Failover Peer IP," and you tell Device B to look for Device A at its respective "Failover Peer IP."

NEW QUESTION: 17
Which statement is true concerning cookie persistence?
A. Cookie persistence allows persistence independent of IP addresses.
B. Cookie persistence allows persistence even if the data are encrypted from client to pool member.
C. Cookie persistence uses a cookie that stores the virtual server, pool name, and member IP address in clear text.
D. If a client's browser accepts cookies, cookie persistence will always cause a cookie to be written to the client's file system.
Answer: (SHOW ANSWER)
Cookie Persistence is a Layer 7 persistence method that leverages an HTTP cookie to track a user session.
* IP Independence: Unlike "Source Address Affinity" (which relies on the client's IP), Cookie persistence identifies the session base16d on a unique token provided by the BIG-IP system. This is crucial for environments where many users share a single gateway (NAT) or where a client's IP might change mid-session.
* Encryption and Decryption: For the BIG-IP to insert or read a cookie, it must be able to see the HTTP header. If the traffic is encrypted end-to-end (SSL Pass-through), the BIG-IP cannot use cookie persistence. SSL must be terminated at the BIG-IP (Option B is false).
* Security: By default, BIG-IP cookies are encoded, not clear text. Modern versions allow for easy encryption of these cookies to prevent information leakage (Option C is false).
* Memory vs. Disk: The default behavior is "session-based" (In-memory). A cookie is only written to the client's file system (disk) if an Expiration is configured in the persistence profile (Option D is false).

NEW QUESTION: 18
What type of virtual server has a destination of 0.0.0.0 and listens on a specific VLAN? (Choose one answer)
A. Standard
B. Forwarding (Layer 2)
C. Wildcard
D. Forwarding (IP)
Answer: (SHOW ANSWER)

In the F5 BIG-IP system, virtual servers are categorized based on their destination address and mask. The system distinguishes between three primary destination scopes:
* Host Virtual Server: A virtual server that has a specific IP address (e.g., 10.10.10.50) and a /32 mask.
* Network Virtual Server: A virtual server that has a destination address representing a subnet (e.g.,
192.168.10.0) and a specific mask (e.g., /24).
* Wildcard Virtual Server: A virtual server that has a destination address of 0.0.0.0 (or :: for IPv6) and a mask of 0.0.0.0 (or /0).
While a "Forwarding (IP)" virtual server (Option D) is the Type (behavioral configuration) often used to route traffic without load balancing, the term Wildcard (Option C) is the specific administrative term used to define the "type" of virtual server based on the 0.0.0.0 destination address.
A common architectural use case is to create a Wildcard Virtual Server that listens only on an internal VLAN to act as a default gateway for outbound traffic (Internet access) for back-end servers. This ensures the BIG-IP system can process and forward traffic that does not match any other specific virtual server configuration.

NEW QUESTION: 19
Which event is always triggered when a client initially connects to a virtual server configured with an HTTP profile?
A. HTTP_DATA
B. CLIENT_DATA
C. HTTP_REQUEST
D. CLIENT_ACCEPTED
Answer: (SHOW ANSWER)
The BIG-IP processing flow follows a specific sequence of events as a packet moves through the system.
* TCP Handshake: Before any application-layer data (like HTTP) can be processed, a TCP connection must be established.
* The First Event: The very first event triggered when a client completes the 3-way handshake with the virtual server is CLIENT_ACCEPTED.
* Profile Influence: Even if an HTTP profile is attached, the system must first "accept" the connection at the protocol level. HTTP_REQUEST only triggers after the client sends data that the BIG-IP recognizes as a valid HTTP request. If a client connects but never sends a request, CLIENT_ACCEPTED will have fired, but HTTP_REQUEST will not.