F5.F5CAB2.v2026-05-19.q27

Exam Code:	F5CAB2
Exam Name:	BIG-IP Administration Data Plane Concepts (F5CAB2)
Certification Provider:	F5
Free Question Number:	27
Version:	v2026-05-19
# of views:	102
# of Questions views:	272
https://www.freecram.net/torrent/F5.F5CAB2.v2026-05-19.q27.html	

NEW QUESTION: 1
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: 2
A BIG-IP Administrator configuresremote authenticationand needs to ensure that users can still log ineven when the remote authentication server is unavailable. Which action should the BIG-IP Administrator take in the remote authentication configuration to meet this requirement? (Choose one answer)
A. Configure a second remote user directory
B. Configure a remote role group
C. Enable theFallback to Localoption
D. Set partition access toAll
Answer: (SHOW ANSWER)
Although remote authentication (LDAP, RADIUS, TACACS+) is acontrol-plane / management-plane feature, it directly affectsavailability and resiliency of administrative access, which is a critical operational HA consideration.
How BIG-IP Remote Authentication Works:
* BIG-IP can authenticate administrators against:
* LDAP
* RADIUS
* TACACS+
* When remote authentication is enabled, BIG-IPby default relies on the remote serverfor user authentication
* If the remote authentication server becomes unreachable, administrators may belocked outunless fallback is configured Why "Fallback to Local" Is Required:
TheFallback to Localoption allows BIG-IP to:
* Attempt authentication against theremote authentication server first

* If the remote server isunreachable or unavailable, fall back to:
* Local BIG-IP user accounts (admin, or other locally defined users)
This ensures:
* Continuous administrative access
* Safe recovery during:
* Network outages
* Authentication server failures
* Maintenance windows
This behavior is explicitly recommended as abest practicein BIG-IP administration to avoid loss of management access.
Why the Other Options Are Incorrect:
* A. Configure a second remote user directory
* Provides redundancyonly if both directories are reachable
* Does not help if remote authentication as a whole is unavailable
* B. Configure a remote role group
* Maps remote users to BIG-IP roles
* Does not affect authentication availability
* D. Set partition access to "All"
* Controls authorization scope after login
* Has no impact on authentication success
Key Availability Concept Reinforced:
To maintainadministrative access resiliency, BIG-IP administrators should always enableFallback to Local when using remote authentication. This prevents lockouts and ensures access even during authentication infrastructure failures.

NEW QUESTION: 3
A BIG-IP Administrator has a cluster of devices. What should the administrator do after creating a new Virtual Server on device 1?
A. create a new cluster on device 1
B. create the new virtual server on device 2
C. synchronize the settings of the group to device 1
D. synchronize the settings of device 1 to the group
Answer: (SHOW ANSWER)
F5 BIG-IP uses aConfigSyncmechanism to ensure that all members of a Device Service Cluster (DSC) share the same configuration.
* Manual Synchronization:By default, configuration changes made on one device (the "source") do not automatically propagate to other members.
* Direction of Sync:Once a Virtual Server is created on device 1, that device's configuration is now
"newer" than the rest of the group. The administrator must initiate a synchronization from themodified device (device 1) to the Sync-Failover group.

* Consistency:This ensures that if a failover occurs, device 2 (the standby) will have the exact same Virtual Server configuration and can take over traffic immediately without interruption.

NEW QUESTION: 4
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:TheFailover Peer IPis the IP address belonging to theotherBIG-IP device in the HA pair.
This is typically a34Self-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: 5
Refer to the exhibit.

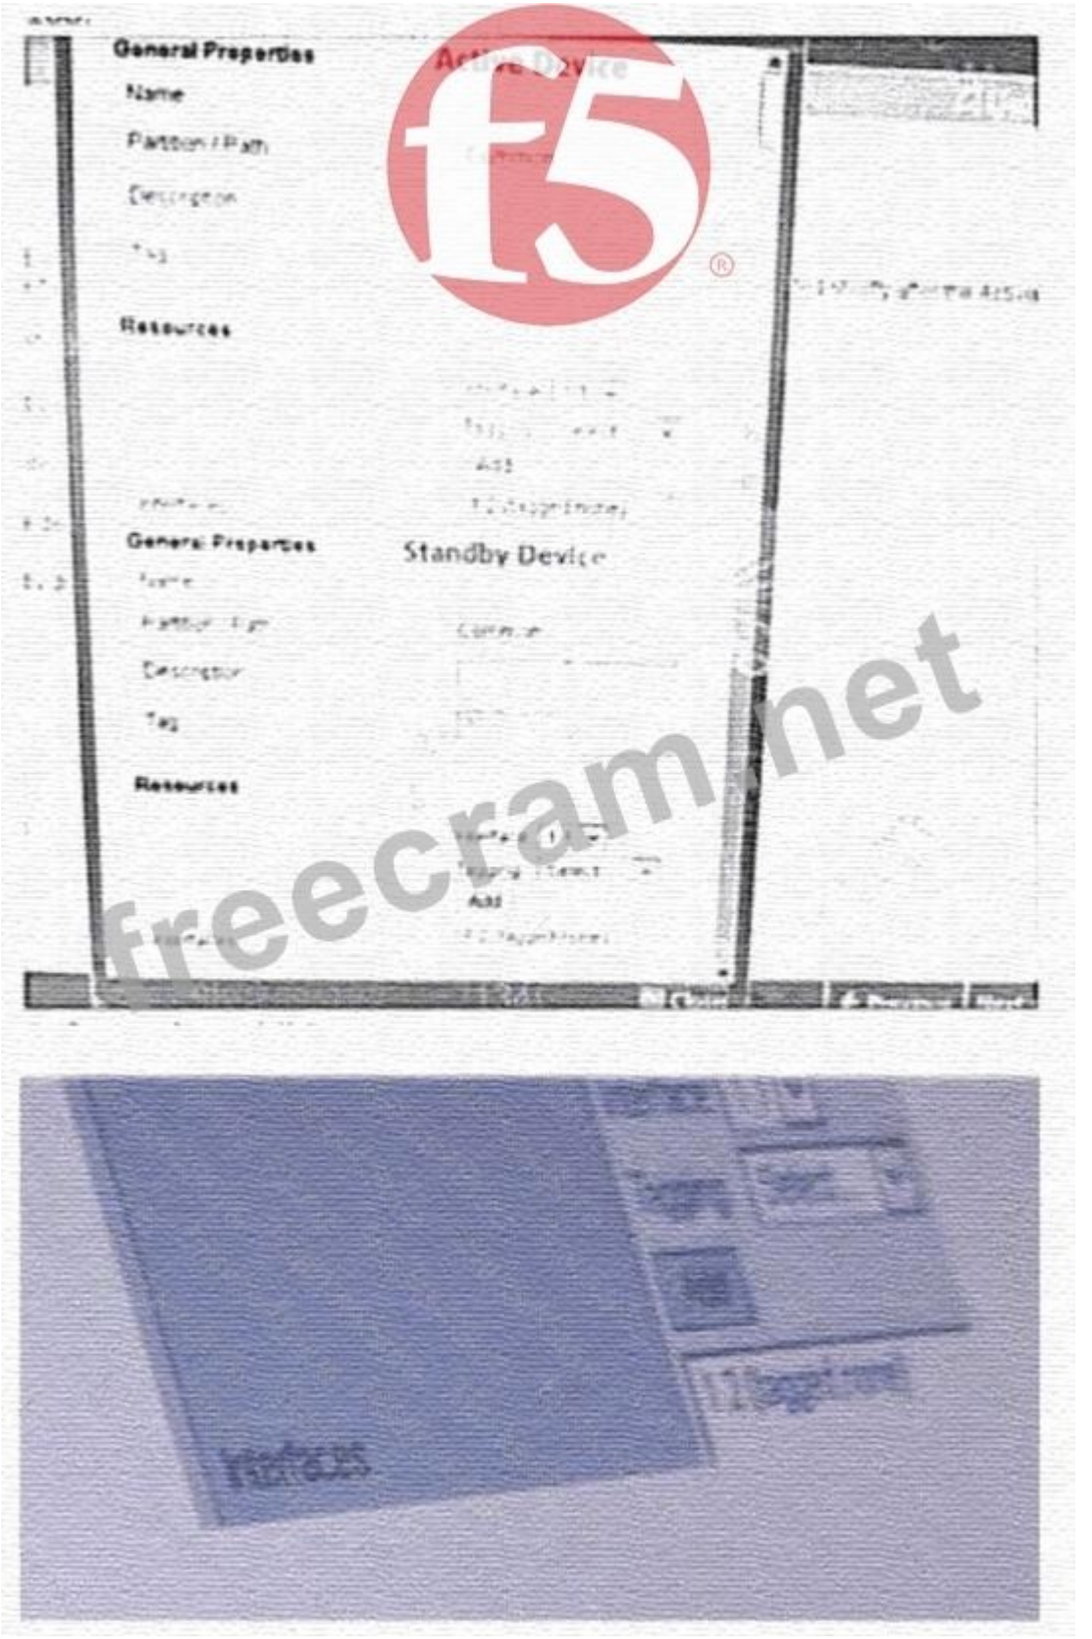

During a planned upgrade to a BIG-IP HA pair running Active/Standby, an outage to application traffic is reported shortly after the Active unit is forced to Standby. Reverting the failover resolves the outage. What should the BIG-IP Administrator modify to avoid an outage during the next failover event?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

A. The Tag value on the Standby device

B. The interface on the Active device to 1.1

C. The Tag value on the Active device
D. The Interface on the Standby device to 1.1
Answer: (SHOW ANSWER)
In anActive/StandbyBIG-IP design, application availability during failover depends on both units having equivalent data-plane connectivityfor the networks that carry application traffic.
Specifically:
* VLANs are bound to specific interfaces (and optionally VLAN tags).
* Floating self IPs / traffic groupsmove to the new Active device during failover.
* For traffic to continue flowing after failover, the new Active device must have thesame VLANs available on the correct interfacesthat connect to the upstream/downstream networks.
What the symptom tells you:
* Traffic works when Device A is Active
* Traffic fails when Device B becomes Active
* Failback immediately restores traffic
This pattern strongly indicates theStandby unit does not have the VLAN connected the same way(wrong physical interface assignment), so when it becomes Active, it owns the floating addresses but cannot actually pass traffic on the correct network segment.
WhyInterface mismatchis the best match:
* If theActiveunit is already working, its interface mapping is correct.
* The fix is to make theStandbyunit's VLAN/interface assignment match the Active unit.
* That corresponds tochanging the Standby device interface to 1.1.
Why theTagoptions are less likely here (given the choices and the exhibit intent):
* Tag issues can also break failover traffic, but the question/options are clearly driving toward the classic HA requirement:consistent VLAN-to-interface mapping on both devicesso the data plane remains functional after the traffic group moves.
Conclusion:To avoid an outage on the next failover, the BIG-IP Administrator must ensure the Standby device uses thesame interface (1.1)for the relevant VLAN(s) that carry the application traffic, so when it becomes Active it can forward/receive traffic normally.

NEW QUESTION: 6
A BIG-IP Administrator is conducting maintenance on one BIG-IP appliance in an HA Pair. Why should the BIG-IP Administrator put the appliance into FORCED-OFFLINE state?
A. To terminate connections to the management IP and decrease persistent connections
B. To allow new connections to Virtual Servers and ensure the appliance becomes active
C. To preserve existing connections to Virtual Servers and reduce the CPU load
D. To terminate existing connections to Virtual Servers and prevent the appliance from becoming active
Answer: (SHOW ANSWER)
TheForced Offlinestate is a critical administrative tool used during maintenance to ensure a device remains in a non-functional state relative to the traffic group.

* Preventing Active Status:When a device is in "Forced Offline," it is effectively disqualified from the HA election process. Even if the other peer fails, a device in Forced Offline willnotbecome active. This is vital during maintenance (like firmware upgrades or hardware replacement) to prevent an unstable or half-configured device from attempting to process traffic.
* Traffic Termination:Placing a device in Forced Offline triggers the system to stop accepting new connections and, depending on the configuration, can facilitate the termination of existing connections so that the administrator can perform work without the data plane actively utilizing system resources.
* Persistence Handling:Unlike the "Disabled" state, Forced Offline ignores persistence records, ensuring that no new traffic is steered to the device via session affinity.

NEW QUESTION: 7
Which statement is true concerning the default communication between a redundant pair of BIG-IP devices?
A. Communication between the systems cannot be effected by port lockdown settings.
B. Data for both connection and persistence mirroring are shared through the same TCP connection.
C. Regardless of the configuration, some data is communicated between the systems at regular intervals.
D. Connection mirroring data is shared through the serial fail over cable unless network failover is enabled.
Answer: (SHOW ANSWER)
Redundant BIG-IP systems (HA pairs) must maintain constant communication to monitor the health of the peer and synchr15onize states.16
* Heartbeats:By default, even with a serial cable, th17e BIG-IP systems exchange "heartbeat" packets over the network to determine if the peer is still alive.
* Network Failover:This involves the exchange of UDP packets (typically on port 1026) at regular intervals.
* Device Service Clustering (DSC):Modern BIG-IP versions use the Central Management (cm) infrastructure to communicate configuration status and sync status constantly.
* Clarification on others:Port lockdowndoesaffect HA communication if misconfigured (A is false). Mirroring uses separate channels (B is false). Mirroring isneversent over the serial cable because it requires high bandwidth (D is false).

NEW QUESTION: 8
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: 9
A BIG-IP Administrator needs to have a BIG-IP linked to two upstream switches for resilience of the external network. The network engineer who is going to configure the switch instructs the BIG-IP Administrator to configure interface binding with LACP. Which configuration should the administrator use?
A. A Trunk containing an interface connected to each switch
B. A Trunk listing the allowed VLAN IDs and MAC addresses configured on the switches
C. A virtual server with an LACP profile and the interfaces connected to the switches as pool members
D. A virtual server with an LACP profile and the switches' management IPs as pool members
Answer: A (LEAVE A REPLY)
In BIG-IP terminology, aTrunkis the object used to implement Link Aggregation (IEEE 802.3ad/802.1AX).

When a network engineer refers to "interface binding" or "EtherChannel" with LACP, the BIG-IP equivalent is a Trunk.
* LACP (Link Aggregation Control Protocol):This is a protocol that allows the BIG-IP system to communicate with the upstream switches to negotiate the bundling of multiple physical links into a single logical link.
* Resilience and Redundancy:By creating a trunk that includes interfaces connected to two different switches (typically configured as a VPC, VSS, or MLAG cluster on the switch side), the administrator ensures that the BIG-IP remains reachable even if one physical interface or one switch fails.
* Data Plane Logic:The BIG-IP treats the trunk as a single Layer 2 interface. VLANs are then associated with the trunk rather than individual physical ports.
Why the other options are incorrect:
* Option B:Trunks aggregate physical interfaces. While VLANs are associated with trunks, the trunk configuration itself does not "list" MAC addresses of the switches; it uses LACP to negotiate the connection.
* Options C & D:Virtual Servers are Layer 4-7 objects used for traffic processing and load balancing.
They do not possess "LACP profiles," nor are physical interfaces or management IPs treated as pool members for the purpose of link aggregation.

NEW QUESTION: 10

Active connections to pool members are unevenly distributed. The load balancing method isLeast Connections (member).Priority Group Activation is disabled.
What is a potential cause of the uneven distribution? (Choose one answer)

A. Priority Group Activation is disabled
B. SSL Profile Server is applied
C. A persistence profile is applied
D. Incorrect load balancing method
Answer: (SHOW ANSWER)
WithLeast Connections (member), BIG-IP attempts to send new connections to the pool member with the fewest current connections. In a perfectly "stateless" scenario (no affinity), this often trends toward a fairly even distribution over time.
However,persistence overrides load balancing:
* When apersistence profileis applied, BIG-IP will continue sending a client (or client group) to thesame pool memberbased on the persistence record (cookie / source address / SSL session ID, etc.).
* This means even if another pool member has fewer connections, BIG-IP may still select the persisted member to honor session affinity.
* The result can beuneven active connection counts, even though the configured load balancing method is Least Connections.
Why the other options are not the best cause:

* A. Priority Group Activation is disabledPriority Group Activation only affects selection when priority groups are configured; disabling it does not inherently create uneven distribution under Least Connections.
* B. SSL Profile Server is appliedA server-side SSL profile affects encryption to pool members, but it does not by itself cause skewed selection across pool members. (Skew could happen indirectly if members have different performance/latency, but that's not the primary, expected exam answer.)
* D. Incorrect load balancing methodLeast Connections is a valid method and does not itself explain unevenness unless something is overriding it (like persistence) or pool members are not all eligible.
Conclusion:
Apersistence profileis the most common and expected reason that active connections becomeunevenly distributed, because persistencetakes precedence overthe Least Connections load-balancing decision.

NEW QUESTION: 11
A standard virtual server has been associated with a pool with multiple members. Assuming all other settings are left at their defaults, which statement is always true concerning traffic processed by the virtual server?
A. The client IP address is unchanged between the client-side connection and the server-side connection.
B. The server IP address is unchanged between the client-side connection and the server-side connection.
C. The TCP ports used in the client-side connection are the same as the TCP ports server-side connection.
D. The IP addresses used in the client-side connection are the same as the IP addresses used in the server- side connection.
Answer: (SHOW ANSWER)
Understanding the default behavior of aStandard Virtual Serverregarding address and port translation is fundamental to BIG-IP administration.
* Source Address Translation (SNAT):By default, the BIG-IP system doesnotperform Source Address Translation (SNAT). This means that the packet's source IP address (theClient IP) remains preserved as it passes through the BIG-IP to the pool member. This is critical for backend servers to identify the original client for logging and security purposes. Therefore, the client IP address is unchanged between the client-side and server-side connections.
* Destination Address Translation (DAT):By default, a Standard Virtual Serveralwaysperforms Destination Address Translation. The BIG-IP system changes the destination IP from the Virtual Server's IP address to the IP address of the specificPool Memberselected by the load balancing algorithm. Consequently, the server-side destination IP is different from the client-side destination IP.

* Port Translation:By default,Port Translationis enabled. If a Virtual Server is listening on port 80 and the selected pool member is configured for port 8080, the BIG-IP will translate the destination port.
Even if the ports happen to be the same, thesettingallows for change, whereas the default SNAT setting (None) ensures the client IP remains static.

NEW QUESTION: 12
An application is configured so that thesame pool member must be used for an entire session, and this behavior must persist acrossHTTP and FTP traffic. A user reports that a session terminates and must be restarted after theactive BIG-IP device fails overto the standby device. Which configuration settings should the BIG-IP Administrator verify to ensure proper behavior when BIG-IP failover occurs? (Choose one answer)
A. Cookie persistence and session timeout
B. Stateful failover and Network Failover detection
C. Persistence mirroring and Match Across Services
D. SYN-cookie insertion threshold and connection low-water mark
Answer: (SHOW ANSWER)
This scenario combinessession continuity,multiple protocols (HTTP and FTP), andHA failover behavior, which directly implicatespersistence handling across devices and services.
Key Requirements Breakdown
* Samepool member for entire session
* Session must survivefailover
* Session must spanmultiple services (HTTP and FTP)
Why Persistence Mirroring + Match Across Services Is Required
Persistence Mirroring
* Ensurespersistence records are synchronizedfrom the active BIG-IP to the standby BIG-IP.
* Without mirroring:
* After failover, the standby device hasno persistence table
* Clients are load-balanced again
* Sessions break, forcing users to restart
* Persistence mirroring is essential forsession continuity during failover Match Across Services
* Allows asingle persistence recordto be shared acrossmultiple virtual servers / protocols
* Required when:
* HTTP and FTP must use thesame pool member
* Multiple services are part of a single application session
Together, these settings ensure:
* Persistence survivesdevice failover
* Persistence is honoredacross HTTP and FTP
Why the Other Options Are Incorrect
* A. Cookie persistence and session timeoutCookie persistence only applies to HTTP and does not address FTP or failover synchronization.

* B. Stateful failover and Network Failover detectionStateful failover applies toconnection state, not persistence records, and does not link HTTP and FTP sessions.
* D. SYN-cookie insertion threshold and connection low-water markThese are DoS / SYN flood protection settings, unrelated to persistence or HA behavior.

NEW QUESTION: 13
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: (SHOW ANSWER)
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: 14
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: 15
A BIG-IP is configured with a pool member located on a different subnet that is not local to the BIG-IP. To ensure that the return traffic from the pool member is sent to the client through the BIG-IP, a Source NAT (SNAT) is used and configured for SNAT Automap. The BIG-IP has a default gateway on the external VLAN, a floating and non-floating self-IP address on each VLAN, and a management address. Which IP address will the BIG-IP use as the source address for the traffic to the pool member when client traffic is sent through the virtual server?
A. The source address will be the first address available in the list of self-IPs.
B. The source address will be the floating self-IP address on the egress VLAN.
C. The source address will be the non-floating self-IP address on the egress VLAN.
D. The source address will be the management IP address.
Answer: (SHOW ANSWER)
SNAT Automapis a feature that automatically selects a self-IP address to use as the source address for translated packets. The selection logic follows a strict hierarchy to ensure that traffic is routable back to the BIG-IP:
* Egress VLAN Priority:The BIG-IP first looks at the VLAN through which the traffic is exiting toward the pool member (the egress VLAN).
* Floating Self-IP Preference:If the egress VLAN has afloating self-IP address, the BIG-IP will always prefer it for SNAT Automap. This is critical for High Availability (HA) because, during a failover, the floating IP moves to the new active device, allowing existing connections to be maintained or correctly timed out.
* Non-Floating Fallback:If no floating self-IP is available on the egress VLAN, the system will use a floating self-IP from a different VLAN. If no floating IPs exist at all, it will then fall back to the non- floating self-IP.
Key Data Plane Concept:
The management IP is never used for data plane traffic. In this scenario, since the administrator has configured a floating self-IP, that specific address becomes the source for all SNAT Automap traffic leaving that VLAN to ensure symmetric routing during HA events.

NEW QUESTION: 16
What command will assist the BIG-IP Administrator in finding the tmm routes when in the TMSH CLI?
A. list net route
B. show net route
C. list net
D. show net
Answer: (SHOW ANSWER)
In the Traffic Management Shell (TMSH), there is a distinct difference between list and show commands.

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

NEW QUESTION: 17
Which statement is true concerning iRule events?
A. All iRule events relate to HTTP processes.
B. All client traffic has data that could be used to trigger iRule events.
C. All iRule events are appropriate at any point in the client-server communication.
D. If an iRule references an event that doesn't occur during the client's communication, the client's connection will be terminated prematurely.
Answer: (SHOW ANSWER)
iRules are event-driven scripts that allow for advanced traffic manipulation.
* Universality of Events:Every packet that passes through t21he BIG-IP data plane triggers events. Even non-HTTP traffic triggers events such as CLIENT_ACCEPTED (when the TCP connection is established22) or CLIENT_DATA (when raw data is received). Therefore, all client traffic-regardless of protocol-has data that can trigger an iRule event.
* Event Specificity:Events are not universal (Option C is false). For example, HTTP_REQUEST only occurs after a full HTTP header is parsed. You cannot trigger an HTTP_RESPONSE event before a request has been sent to a server.
* Protocol Agnostic:iRules are not limited to HTTP (Option A is false); they can handle TCP, UDP, DNS, FTP, SIP, and more.
* Error Handling:If an iRule references an event that never triggers (e.g., an HTTP_REQUEST event in a purely TCP virtual server), the iRule code for that event simply never executes. It doesnotterminate the connection (Option D is false).

NEW QUESTION: 18
Which of the following lists theorder of preference from most preferred to least preferredwhen BIG-IP processes and selects avirtual server? (Choose one answer)
A. Destination host address # Source host address # Service port
B. Source host address # Service port # Destination host address
C. Service port # Destination host address # Source host address
Answer: (SHOW ANSWER)
The BIG-IP system uses a specific precedence algorithm to determine which virtual server (listener) should process an incoming packet when multiple virtual servers might match the criteria. Since BIG-IP version
11.3.0, the system evaluates three primary factors in a fixed order of importance:
* Destination Address:The system first looks for the most specific destination match. A "Host" address (mask /32) is preferred over a "Network" address (mask /24, /16, etc.), which is preferred over a
"Wildcard" (0.0.0.0/0).
* Source Address:If multiple virtual servers have identical destination masks, the system then evaluates the source address criteria. Again, a specific source host match is preferred over a source network or a wildcard source.
* Service Port:Finally, if both destination and source specifications are equal, the system checks the port.
A specific port match (e.g., 80) is preferred over a wildcard port (e.g., or 0).
Following this logic, a virtual server configured with a specificdestination host, a specificsource host, and a specificservice portrepresents the highest level of specificity and thus the highest preference.

NEW QUESTION: 19
The network architecture for a BIG-IP consists of anexternal VLANand aninternal VLANwithtwo interfaces connected to the upstream switch. The design requiresfault tolerancein the case that one of the interfaces is down. Which deployment architecture meets these requirements? (Choose one answer)
A. One network trunk with both VLANs and LACP enabled, and both VLANs configured asuntagged
B. Two network trunks each with one VLAN and LACP enabled, and both VLANs configured astagged
C. Two network trunks each with one VLAN and LACP disabled, and one VLAN configured astaggedand one VLAN configured asuntagged
D. One network trunk with both VLANs and LACP enabled, and both VLANs configured astagged
Answer: (SHOW ANSWER)
To meet the requirement offault tolerance when one interface goes down, BIG-IP must uselink aggregation so that loss of a single physical link does not isolate the VLAN(s).

How the objects relate (data plane view)
* Interfaces= physical links.
* Trunk (LACP)= bundles multiple interfaces into one logical link that providesredundancy(and possibly bandwidth aggregation).
* VLANsare assigned to interfaces or trunks. If you needmultiple VLANs on the same trunk, they must use802.1Q tagging(because you can only have one untagged VLAN per interface/trunk).
* Self IPsare then placed on the VLANs to provide BIG-IP presence and routing/ARP functions, but self IPs are not what provides link resiliency-the trunk does.
Why Option D is correct
* You havetwo physical interfacesand you want resiliency if one fails # put both interfaces intoone trunk with LACP enabled.
* You needboth external and internal VLANson those same two links # both VLANs should be configured astaggedon that trunk, so they can coexist on the same aggregated link.
* If either physical interface fails, the trunk remains up via the remaining interface, keepingboth VLANs operational.
Why the other options are incorrect
* A: Two VLANs cannot both beuntaggedon the same trunk/interface. Only one untagged VLAN is possible; additional VLANs must be tagged.
* B: Two trunks "each with one VLAN" would typically mean splitting VLANs across separate trunks.
With only two interfaces total, that becomes one interface per trunk-if one interface goes down, the VLAN on that interface is down (no redundancy for that VLAN).
* C: Same redundancy problem as B, and disabling LACP removes the negotiated aggregation behavior expected when the switch engineer specifically requested LACP.

NEW QUESTION: 20
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: 21
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 > Nodes > Default Monitor
C. Local Traffic > Pools > DB_Pool > Members
D. Local Traffic > Pools > DB_Pool > Properties
Answer: (SHOW ANSWER)
In the BIG-IP system object hierarchy, health monitors can be applied at three levels:Node,Pool, andPool Member.
* Pool Level (Properties):Applying a monitor at thePool > Propertieslevel is the most common and efficient administrative practice. When applied here, the monitor is inherited by all members of that pool. If the monitor fails for a specific member, that member is marked "down" specifically for that pool.
* Node Level:If a monitor is applied at the Node level (Local Traffic > Nodes), it checks the health of the physical IP address itself. If it fails, that node (and all pool members associated with it) is marked down globally across the entire system.
* Member Level:Applying a monitor at thePool > Memberslevel allows for specific "per-member" monitoring, which is usually only done if different members in the same pool require different health checks.
* The Specific Case:For a standard database pool like DB_Pool, the administrator should navigate to Local Traffic > Pools > DB_Pool > Propertiesand select the appropriate monitor (e.g., a custom TCP or Oracle monitor) from the "Health Monitors" configuration section.

NEW QUESTION: 22
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: 23

A BIG-IP Administrator is informed that traffic oninterface 1.1is expected to increase beyond the maximum bandwidth capacity of the link. There is asingle VLANon the interface.

What should the BIG-IP Administrator do to increase the total available bandwidth? (Choose one answer)

A. Increase the MTU on the VLAN using interface 1.1
B. Create a trunk object with two interfaces
C. Assign two interfaces to the VLAN
D. Set the media speed of interface 1.1 manually

Answer: (SHOW ANSWER)

On BIG-IP systems,physical interface bandwidth is fixedby the link speed (for example, 1GbE or 10GbE).

When traffic demand exceeds the capacity of a single interface, BIG-IP provideslink aggregationthrough trunks.

Key concepts involved:

* InterfacesA single physical interface (such as 1.1) is limited to its negotiated link speed. You cannot exceed this capacity through software tuning alone.

* Trunks (Link Aggregation)Atrunkcombines multiple physical interfaces into a single logical interface.
* BIG-IP supports LACP and static trunks.
* Traffic is distributed across member interfaces, increasingaggregate bandwidthand providing redundancy.
* VLANs are then assigned to thetrunk, not directly to individual interfaces.
Why optionBis correct:
* Creating atrunk with two interfacesallows BIG-IP to use both physical links simultaneously.
* This increases total available bandwidth (for example, two 10Gb interfaces # up to 20Gb aggregate capacity).
* This is thedocumented and supported methodfor scaling bandwidth on BIG-IP.
Why the other options are incorrect:
* A. Increase the MTUMTU changes affect packet size and efficiency,not total bandwidth capacity.
* C. Assign two interfaces to the VLANBIG-IP doesnotsupport assigning a VLAN to multiple interfaces directly. VLANs must be associated withone interface or one trunk.
* D. Set the media speed manuallyMedia speed can only be set up to the physical capability of the interface and connected switch port. It cannot exceed the hardware limit.
Conclusion:
To increase total available bandwidth on BIG-IP when a single interface is insufficient, the administrator must create atrunk object with multiple interfacesand move the VLAN onto the trunk. This aligns directly with BIG-IP data plane design and best practices.

NEW QUESTION: 24
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: 25
What type of Virtual Server is configured with no Pool-members, and proxies traffic to the destination IP address specified by the client device?
A. Forwarding (IP)
B. Standard
C. Performance (Layer 4)
D. Stateless
Answer: (SHOW ANSWER)
AForwarding (IP)virtual server is unique because it does not perform load balancing in the traditional sense.
* No Pool Members:Unlike a Standard virtual server, which requires a pool to direct traffic, a Forwarding (IP) virtual server typically hasno pool assigned.
* Destination-Based Routing:The BIG-IP system looks at thedestination IP addressin the original packet header sent by the client. It then consults the BIG-IP system's local routing table to determine where to send the packet.
* Transparency:It acts as a high-performance router/gateway, often used to forward traffic from internal servers to the internet or across different subnets while still allowing the BIG-IP to apply features like SNAT or bandwidth controllers.
* Stateful Tracking:While it forwards traffic based on the routing table, it still creates an entry in the connection table to track the flow (unless it is a Stateless virtual server).

NEW QUESTION: 26
The owner of a web application asks the BIG-IP Administrator to change the port that the BIG-IP device sends traffic to. This change must be made for each member in the server pool named app_pool for the Virtual Server named app_vs. In which area of the BIG-IP Configuration Utility should the BIG-IP Administrator make this change?
A. Local Traffic > Virtual Servers
B. Local Traffic > Pools
C. Local Traffic > Nodes
D. Network > Interfaces
Answer: (SHOW ANSWER)
In the BIG-IP object hierarchy, the destination port for backend traffic is defined at thePool Memberlevel.

While a Virtual Server listens on a specific port, the Pool determines where that traffic is directed after the load balancing decision is made.
* Pools and Pool Members: A pool is a collection of devices, often called pool members, to which the BIG-IP system passes traffic. Each pool member is defined by an IP address and a service port.
* Port Translation: When an administrator needs to change the port the BIG-IP uses to communicate with backend servers, they must navigate to the specific Pool and modify the service port for each member within that pool.
* Logical Separation:
* Virtual Serversdefine the "front-end" port where clients connect.
* Poolsdefine the "back-end" port where the application resides.
* Nodesrepresent the physical server's IP address and do not contain port-specific configuration.

NEW QUESTION: 27
What type of Virtual Server is configured with no Pool-members, and proxies traffic to the destination IP address specified by the client device?
A. Stateless
B. Standard
C. Forwarding (IP)
D. Performance (Layer 4)
Answer: (SHOW ANSWER)