F5.F5CAB2.v2026-07-22.q41

Exam Code:	F5CAB2
Exam Name:	BIG-IP Administration Data Plane Concepts (F5CAB2)
Certification Provider:	F5
Free Question Number:	41
Version:	v2026-07-22
# of views:	125
# of Questions views:	410
https://www.freeqas.com/qa/F5/F5CAB2/F5.F5CAB2.v2026-07-22.q41.html	

NEW QUESTION: 1
The BIG-IP Administrator needs to ensure that if a pool member is marked down by the monitor, the BIG-IP system sends existing connections to another available pool member. Which task should the BIG-IP Administrator perform to meet this goal?
A. Reconfigure the pool monitor to mark the member as UP
B. Enable mirroring within the persistence profile
C. Set Action on Service Down setting under the Virtual Server configuration to reselect
D. Set Action on Service Down setting under the pool configuration to reselect
Answer: D (LEAVE A REPLY)
By default, when a pool member is marked "down" by a monitor, the BIG-IP system stops sendingnew connections to that member, but it typically allows existing connections to time out naturally (or resets them depending on profile settings).
* Action on Service Down:This setting is configured at thePoollevel.
* Reselect:When set toReselect, if a pool member is marked down, the BIG-IP system will immediately attempt to pick a different available pool member for any existing, active connections associated with the failed member.
* Client Experience:This is used to maintain the user session by transparently moving the traffic to a healthy server without the client needing to re-establish the connection to the Virtual Server.

NEW QUESTION: 2
A BIG-IP Administrator is informed that traffic on interface 1.1 is expected to increase beyond the maximum bandwidth capacity of the link. There is a single VLAN on the interface.
What should the BIG-IP Administrator do to increase the total available bandwidth? (Choose one answer)
A. Increase the MTU on the VLAN using interface 1.1
B. Create a trunk object with two interfaces
C. Assign two interfaces to the VLAN
D. Set the media speed of interface 1.1 manually
Answer: B (LEAVE A REPLY)
Comprehensive and Detailed Explanation (BIG-IP Administration - Data Plane Concepts):
On BIG-IP systems, physical interface bandwidth is fixed by the link speed (for example, 1GbE or 10GbE). When traffic demand exceeds the capacity of a single interface, BIG-IP provides link aggregation through trunks.

Key concepts involved:
Interfaces
A single physical interface (such as 1.1) is limited to its negotiated link speed. You cannot exceed this capacity through software tuning alone.
Trunks (Link Aggregation)
A trunk combines multiple physical interfaces into a single logical interface.
BIG-IP supports LACP and static trunks.
Traffic is distributed across member interfaces, increasing aggregate bandwidth and providing redundancy.
VLANs are then assigned to the trunk, not directly to individual interfaces.
Why option B is correct:
Creating a trunk with two interfaces allows BIG-IP to use both physical links simultaneously.
This increases total available bandwidth (for example, two 10Gb interfaces → up to 20Gb aggregate capacity).
This is the documented and supported method for scaling bandwidth on BIG-IP.
Why the other options are incorrect:
A . Increase the MTU
MTU changes affect packet size and efficiency, not total bandwidth capacity.
C . Assign two interfaces to the VLAN
BIG-IP does not support assigning a VLAN to multiple interfaces directly. VLANs must be associated with one interface or one trunk.
D . Set the media speed manually
Media speed can only be set up to the physical capability of the interface and connected switch port. It cannot exceed the hardware limit.
Conclusion:
To increase total available bandwidth on BIG-IP when a single interface is insufficient, the administrator must create a trunk object with multiple interfaces and move the VLAN onto the trunk.
This aligns directly with BIG-IP data plane design and best practices.

NEW QUESTION: 3
What should a BIG-IP Administrator configure to minimize impact during a failover? (Choose one answer)
A. External monitors
B. Clone pool
C. OneConnect profile
D. MAC masquerading
Answer: D (LEAVE A REPLY)
In BIG-IP high availability (HA) deployments, one of the primary causes of traffic disruption during failover is Layer 2 and Layer 3 relearning by upstream network devices (switches and routers). When traffic groups move from the Active device to the Standby device, the network must quickly associate the IP addresses with the new device.
Why MAC Masquerading Minimizes Failover Impact:
MAC masquerading allows a traffic group to use a floating, shared MAC address for its Self IPs. This MAC address moves with the traffic group during failover.
Key benefits:
* The MAC address does not change when failover occurs
* Upstream switches do not need to relearn ARP entries
* Traffic resumes almost immediately after failover
* Dramatically reduces packet loss and connection interruption
From BIG-IP Administration Data Plane Concepts:
* MAC masquerade is specifically designed to provide fast failover

* It is a best practice for HA pairs, especially in environments sensitive to latency and connection loss Why the Other Options Are Incorrect:
* A. External monitors
* Used to check the availability of external resources
* Do not reduce network convergence or failover disruption
* B. Clone pool
* Used for traffic mirroring or security analysis
* Has no impact on failover behavior
* C. OneConnect profile
* Optimizes server-side TCP connections
* Does not address ARP or MAC relearning during failover
Key HA Concept Reinforced:
To minimize failover impact on live traffic, BIG-IP administrators should ensure Layer 2 continuity. MAC masquerading is the primary mechanism that enables near-instant failover by preventing ARP and MAC table reconvergence delays.

NEW QUESTION: 4
An application is configured so that the same pool member must be used for an entire session, and this behavior must persist across HTTP and FTP traffic. A user reports that a session terminates and must be restarted after the active BIG-IP device fails over to 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 combines session continuity, multiple protocols (HTTP and FTP), and HA failover behavior
, which directly implicates persistence handling across devices and services.
Key Requirements Breakdown
* Same pool member for entire session
* Session must survive failover
* Session must span multiple services (HTTP and FTP)
Why Persistence Mirroring + Match Across Services Is Required
Persistence Mirroring
* Ensures persistence records are synchronized from the active BIG-IP to the standby BIG-IP.
* Without mirroring:
* After failover, the standby device has no persistence table
* Clients are load-balanced again
* Sessions break, forcing users to restart
* Persistence mirroring is essential for session continuity during failover Match Across Services
* Allows a single persistence record to be shared across multiple virtual servers / protocols
* Required when:
* HTTP and FTP must use the same pool member
* Multiple services are part of a single application session

Together, these settings ensure:
* Persistence survives device failover
* Persistence is honored across 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 to connection 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: 5
Which of the following lists the order of preference from most preferred to least preferred when BIG-IP processes and selects a virtual 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 specific destination host, a specific source host, and a specific service port represents the highest level of specificity and thus the highest preference.

NEW QUESTION: 6
The BIG-IP Administrator wants to provide quick failover between the F5 LTM devices that are configured as an HA pair with a single Self IP using the MAC Masquerade feature. The administrator configures MAC masquerade for traffic-group-1 using the following command:
`tmsh modify /cm traffic-group traffic-group-1 mac 02:12:34:56:00:00`
However, the Network Operations team identifies an issue with using the same MAC address across multiple VLANs. As a result, the administrator enables Per-VLAN MAC Masquerade to ensure a unique MAC address per VLAN by running:
`tmsh modify /sys db tm.macmasqaddr_per_vlan value true`
What would be the resulting MAC address on a tagged VLAN with ID 1501? (Choose one answer)
A. 02:12:34:56:01:15
B. 02:12:34:56:dd:05
C. 02:12:34:56:15:01
D. 02:12:34:56:05:dd
Answer: D (LEAVE A REPLY)

NEW QUESTION: 7
Refer to the exhibit.

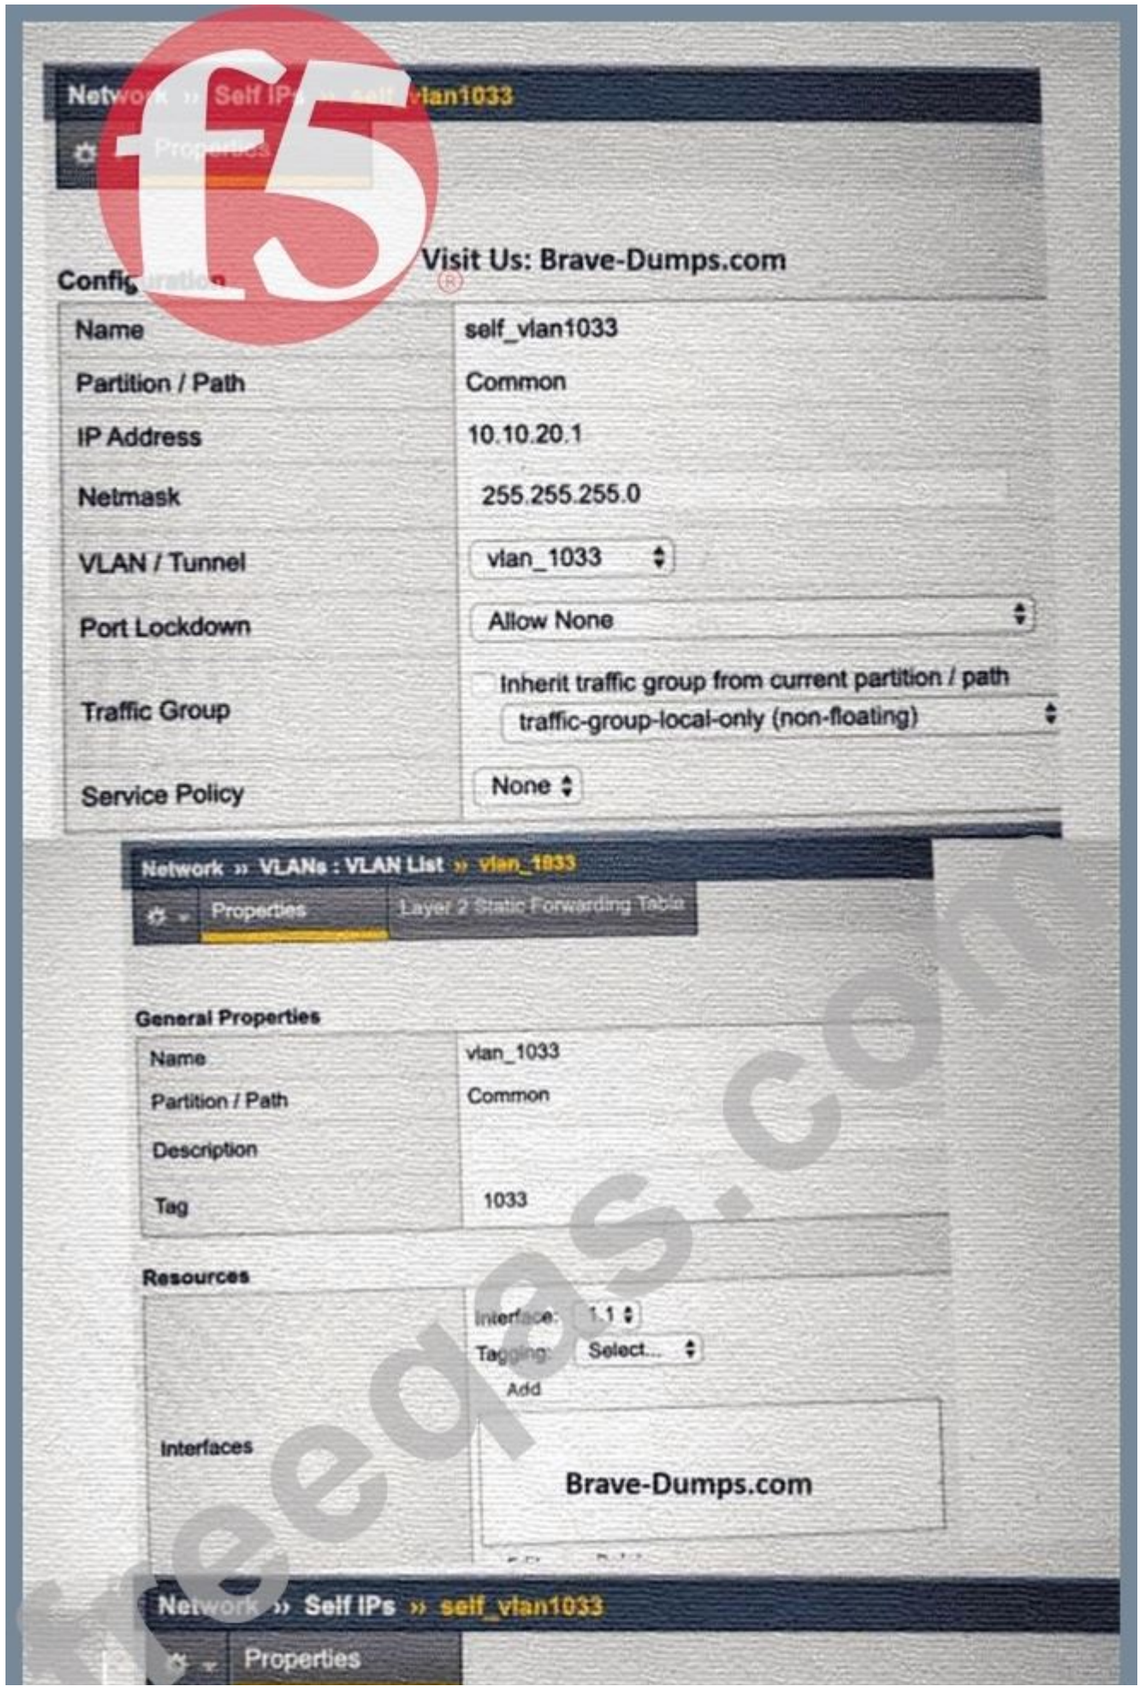
Visit Us: Brave-Dumps.com
Name
self_vlan1033
Partition / Path
Common
IP Address
10.10.20.1
Netmask
255.255.255.0
VLAN / Tunnel
vlan_1033
Port Lockdown
Allow None
Traffic Group
Inherit traffic group from current partition / path
traffic-group-local-only (non-floating)
Service Policy
None
Network » VLANs : VLAN List » vlan_1033
Properties
Layer 2 Static Forwarding Table
General Properties
Name
vlan_1033
Partition / Path
Common
Description
Tag
1033
Resources
Interface:
1.1
Tagging:
Select...
Add
Interfaces
Brave-Dumps.com
Network » Self IPs » self_vlan1033
Properties

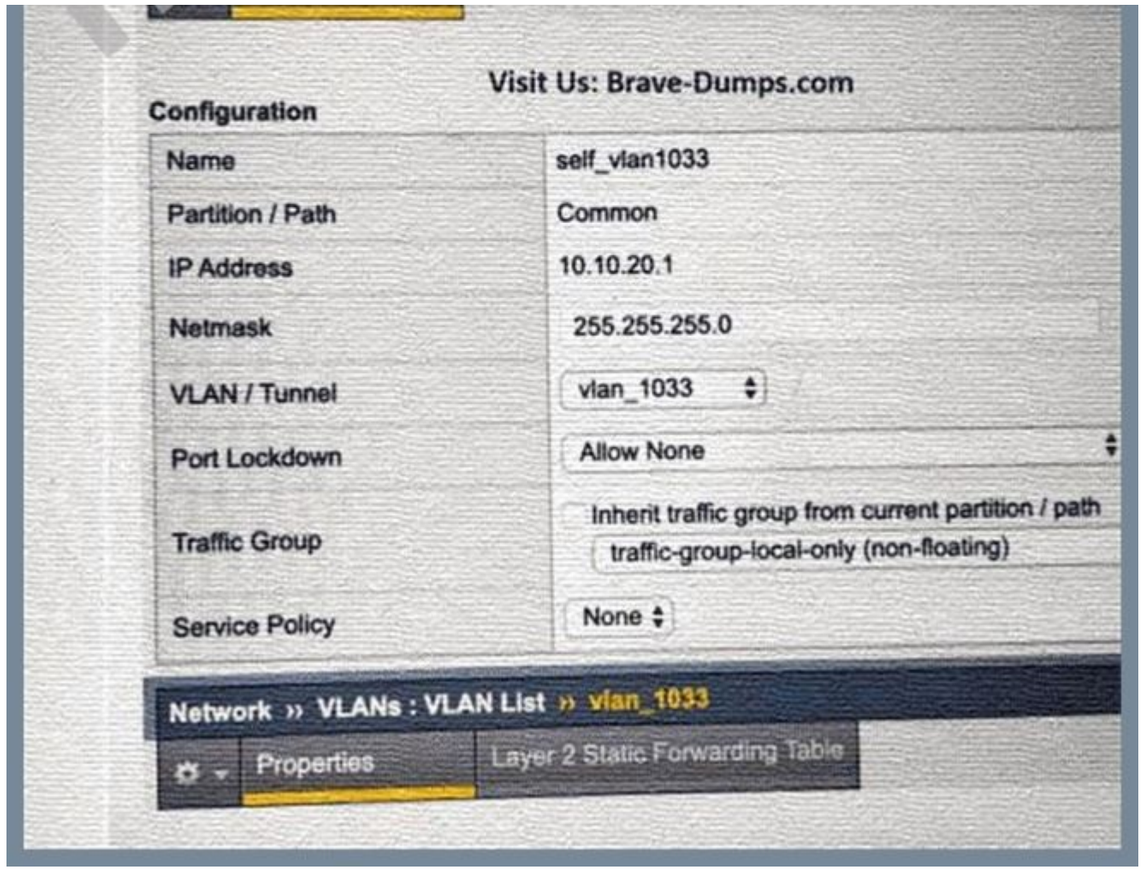

The network team creates a new VLAN on the switches. The BIG-IP Administrator creates a new VLAN and a Self IP on the BIG-IP device, but theservers on the new VLAN are NOT reachablefrom the BIG-IP device.

Which action should the BIG-IP Administrator take to resolve this issue? (Choose one answer)

A. Set Port Lockdown of the Self IP toAllow All

B. ChangeAuto Last Hopto enabled

C. Assign aphysical interfaceto the new VLAN

D. Create aFloating Self IP address

Answer: (SHOW ANSWER)

For BIG-IP to send or receive traffic on a VLAN, that VLANmust be bound to a physical interface or a trunk. Creating a VLAN object and a Self IP alone isnot sufficientto establish data-plane connectivity.

From the exhibit:

* The VLAN (vlan_1033) exists and has a tag defined.
* A Self IP is configured and associated with the VLAN.
* However, traffic cannot reach servers on that VLAN.

This indicates aLayer 2 connectivity issue, not a Layer 3 or HA issue.

Whyassigning a physical interfacefixes the problem:

* BIG-IP VLANs do not carry traffic unless they are explicitly attached to:
* Aphysical interface(e.g., 1.1), or
* Atrunk
* Without an interface assignment, the VLAN is effectively isolated and cannot transmit or receive frames, making servers unreachable regardless of correct IP addressing.

Why the other options are incorrect:
* A. Set Port Lockdown to Allow AllPort Lockdown controls which services can be accessedon the Self IP(management-plane access), not whether BIG-IP can reach servers on that VLAN.
* B. Change Auto Last Hop to enabledAuto Last Hop affects return traffic routing for asymmetric paths. It does not fix missing Layer 2 connectivity.
* D. Create a Floating Self IP addressFloating Self IPs are used for HA failover. They do not resolve reachability issues on a single device when the VLAN itself is not connected to an interface.
Conclusion:
The servers are unreachable because the VLAN hasno physical interface assigned. To restore connectivity, the BIG-IP Administrator mustassign a physical interface (or trunk) to the VLAN, enabling Layer 2 traffic flow.

NEW QUESTION: 8
Refer to the exhibit.

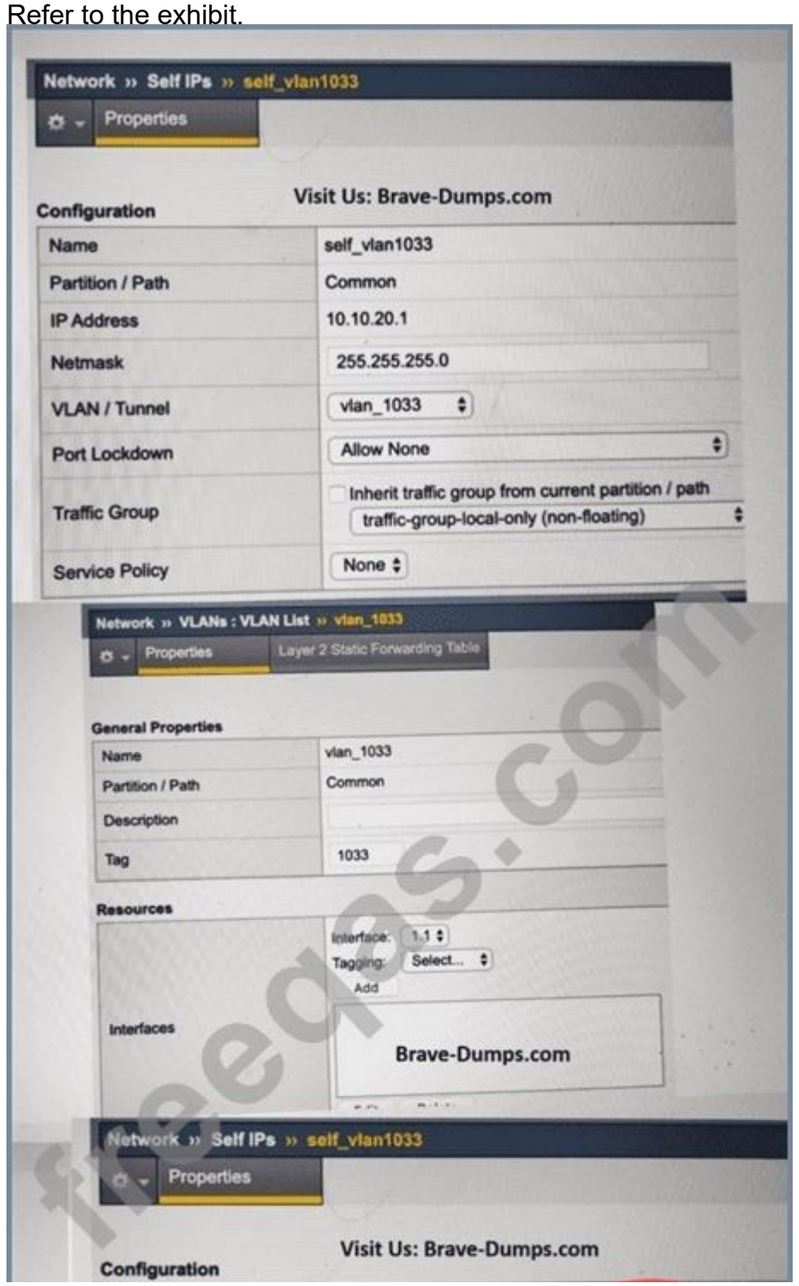

The network team creates a new VLAN on the switches. The BIG-IP Administrator creates a new VLAN and a Self IP on the BIG-IP device, but the servers on the new VLAN are NOT reachable from the BIG-IP device.

Which action should the BIG-IP Administrator take to resolve this issue? (Choose one answer)

A. Set Port Lockdown of the Self IP to Allow All

B. Change Auto Last Hop to enabled

C. Assign a physical interface to the new VLAN

D. Create a Floating Self IP address

Answer: C (LEAVE A REPLY)

For BIG-IP to send or receive traffic on a VLAN, that VLAN must be bound to a physical interface or a trunk. Creating a VLAN object and a Self IP alone is not sufficient to establish data-plane connectivity.

From the exhibit:

* The VLAN (vlan_1033) exists and has a tag defined.
* A Self IP is configured and associated with the VLAN.
* However, traffic cannot reach servers on that VLAN.

This indicates a Layer 2 connectivity issue, not a Layer 3 or HA issue.

Why assigning a physical interface fixes the problem:

* BIG-IP VLANs do not carry traffic unless they are explicitly attached to:
* A physical interface (e.g., 1.1), or
* A trunk
* Without an interface assignment, the VLAN is effectively isolated and cannot transmit or receive frames, making servers unreachable regardless of correct IP addressing.

Why the other options are incorrect:

* A. Set Port Lockdown to Allow AllPort Lockdown controls which services can be accessed on the Self IP (management-plane access), not whether BIG-IP can reach servers on that VLAN.
* B. Change Auto Last Hop to enabledAuto Last Hop affects return traffic routing for asymmetric paths. It does not fix missing Layer 2 connectivity.
* D. Create a Floating Self IP addressFloating Self IPs are used for HA failover. They do not resolve reachability issues on a single device when the VLAN itself is not connected to an interface.

Conclusion:

The servers are unreachable because the VLAN has no physical interface assigned. To restore connectivity, the BIG-IP Administrator must assign a physical interface (or trunk) to the VLAN, enabling Layer 2 traffic flow.

NEW QUESTION: 9
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: (SHOW ANSWER)
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: 10
What is the result when a BIG-IP Administrator manually disables a pool member? (Choose one answer)
A. The disabled pool member stops processing persistent connections.
B. All pool members continue to process persistent connections.
C. The disabled pool member stops processing existing connections.
D. All pool members stop accepting new connections.
Answer: A (LEAVE A REPLY)
Comprehensive and Detailed Explanation From BIG-IP Administration Data Plane Concepts documents:
In BIG-IP LTM, a pool member state directly affects how traffic is handled at the data plane level. When a pool member is manually disabled, BIG-IP changes the member's availability state to disabled, which has specific and predictable traffic-handling consequences.
According to BIG-IP Administration Data Plane Concepts:
A disabled pool member:
Does not accept new connections
Continues to process existing non-persistent connections until they naturally close Is removed from load-balancing decisions, including persistence lookups Most importantly for this question: Persistent connections
(such as those created using source-address persistence, cookie persistence, or SSL persistence) are not honored for a disabled pool member BIG-IP will not send new persistent traffic to a disabled member, even if persistence records exist Therefore, when a pool member is manually disabled, it stops processing persistent connections, while allowing existing non-persistent flows to drain gracefully.
Why the Other Options Are Incorrect:

B - Persistent connections are not honored for a disabled pool member

C - Existing connections are not immediately terminated when a pool member is disabled D - Only the disabled pool member stops accepting new connections, not all pool members Key Data Plane Concept Reinforced:

Manually disabling a pool member is a graceful administrative action that prevents new and persistent traffic from reaching the member while allowing existing connections to complete, which is critical for maintenance and troubleshooting scenarios.

NEW QUESTION: 11

A virtual server is listening at 10.10.1.100:any and has the following iRule associated with it:

when CLIENT_ACCEPTED { if {[TCP::local_port] equals 21 } { pool

ftppool } elseif {[TCP::local_port] equals 23 } { pool telnetpool }

If a user connects to 10.10.1.100 and port 22, which pool will receive the request?

A. ftppool

B. None. The request will be dropped.

C. telnetpool

D. Unknown. The pool cannot be determined from the information provided.

Answer: D (LEAVE A REPLY)

NEW QUESTION: 12

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: 13

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: D (LEAVE A REPLY)

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: 14
Which of the following lists the order of preference from most preferred to least preferred when BIG-IP processes and selects a virtual 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 specific destination host, a specific source host, and a specific service port represents the highest level of specificity and thus the highest preference.

NEW QUESTION: 15
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 > Pools > De Pool > Members
B. Local Traffic > Pools > DB Pool > Properties
C. Local Traffic > Profiles » Protocol > TCP
D. Local Traffic > Nodes > Default Monitor
Answer: (SHOW ANSWER)

NEW QUESTION: 16
What should a BIG-IP Administrator configure to minimize impact during a failover?
A. MAC masquerading
B. External monitors
C. OneConnect profile
D. Clone pool
Answer: A (LEAVE A REPLY)
In a High Availability (HA) environment, a failover event occurs when the active BIG-IP system stops processing traffic and the standby unit takes over. This transition can cause a brief

disruption in network traffic because the surrounding switches need to update their ARP tables to associate the Virtual IP (VIP) and floating Self-IPs with the MAC address of the new active unit.
* MAC Masquerade Functionality:To minimize this impact, an administrator can configureMAC masquerading. This feature allows the administrator to assign a unique, "shared" MAC address to a traffic group.
* Seamless Transition:When a failover occurs, the new active unit begins using this shared MAC address immediately. Since the MAC address for the traffic group remains the same regardless of which physical device is active, the upstream switches do not need to update their ARP tables or learn a new MAC-to-port mapping.
* Packet Loss Reduction:By maintaining a constant MAC address, MAC masquerading significantly reduces the time it takes for traffic to resume, effectively eliminating the "gratuitous ARP" dependency and minimizing packet loss during the handover.
Why other options are incorrect:
* External monitors:These are used for advanced health checking of pool members and do not directly impact the speed or smoothness of a device-level failover.
* OneConnect profile:This is a performance optimization tool that aggregates multiple client-side requests into a single server-side TCP connection; it is not a failover mechanism.
* Clone pool:This is used to replicate traffic for IDS or monitoring purposes and has no role in high availability or failover optimization.

Valid F5CAB2 Dumps shared by PrepPdf.com for Helping Passing F5CAB2 Exam! PrepPdf.com now offer the **newest F5CAB2 exam dumps**, the PrepPdf.com F5CAB2 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com F5CAB2 dumps with Test Engine here: https://www.preppdf.com/F5/F5CAB2-prepaway-exam-dumps.html (**68** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 17
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. Standard
B. Forwarding (IP)
C. Performance (Layer 4)
D. Stateless
Answer: B (LEAVE A REPLY)

NEW QUESTION: 18
Refer to the exhibit above.

Flag for review

Local Traffic » Pools : Pool List » docker_www_farm

Properties | Members | Statistics

Load Balancing

Load Balancing Method	Round Robin		
Priority Group Activation	Less than... 2	Available Member(s)	

Update

Current Members

Add...

Status	Member	Address	Service Port	FQDN	Ephemeral	Ratio	Priority Group	Connection Limit	Partition / Path
	serv1:80	192.168.30.11	80		No	1	2 (Active)	0	Common
	serv2:80	192.168.30.12	80		No	1	2 (Inactive)	0	Common
	serv3:80	192.168.30.13	80		No	1	1 (Active)	0	Common
	serv4:80	192.168.30.14	80		No	1	1 (Active)	0	Common

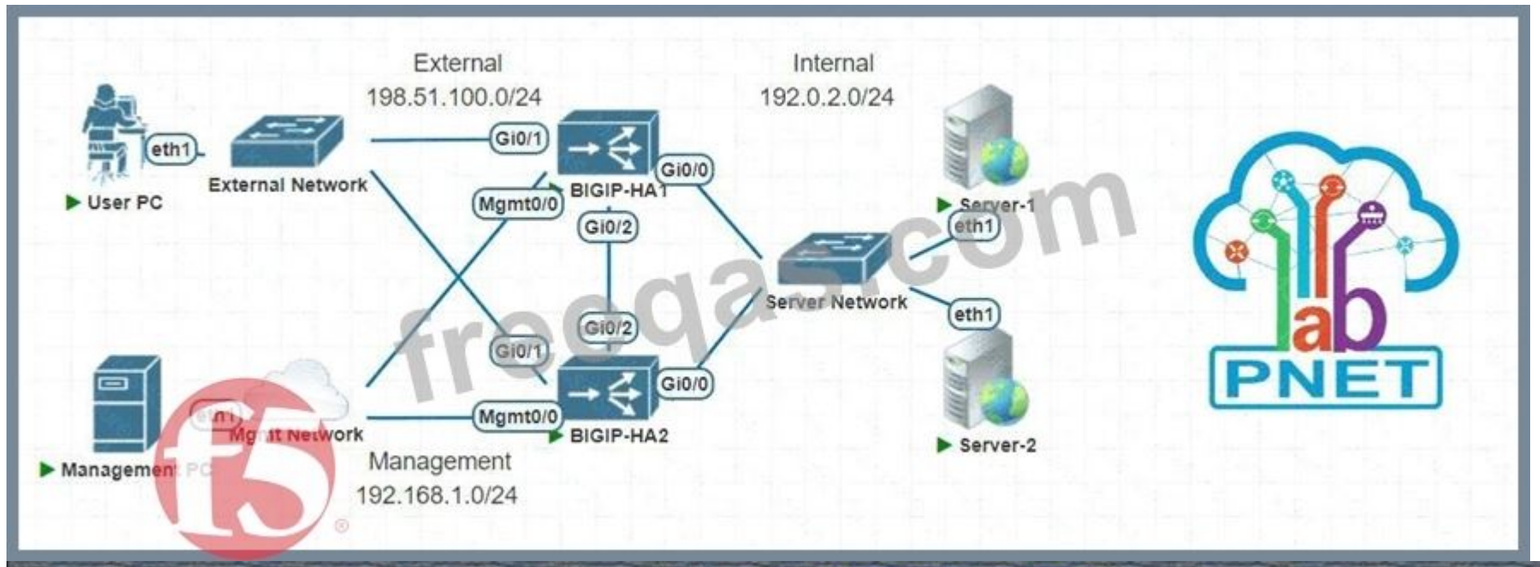

freeqa.com

Local Traffic » Pools : Pool List » www_pool

Properties | Members | Statistics

Member Properties

Node Name	10.1.20.11
Address	10.1.20.11
Service Port	80
Partition / Path	Common
Description	
Parent Node	10.1.20.11
Availability	Available (Enabled) - Pool member is available 2018-05-29 16:56:28

Health Monitors	http
Monitor Logging	Enable
Current Connections	0
State	Enabled (All traffic allowed) Disabled (Only persistent or active connections allowed) Forced Offline (Only active connections allowed)

Configuration: Basic

Ratio	3
Priority Group	0
Connection Limit	0
Connection Rate Limit	0

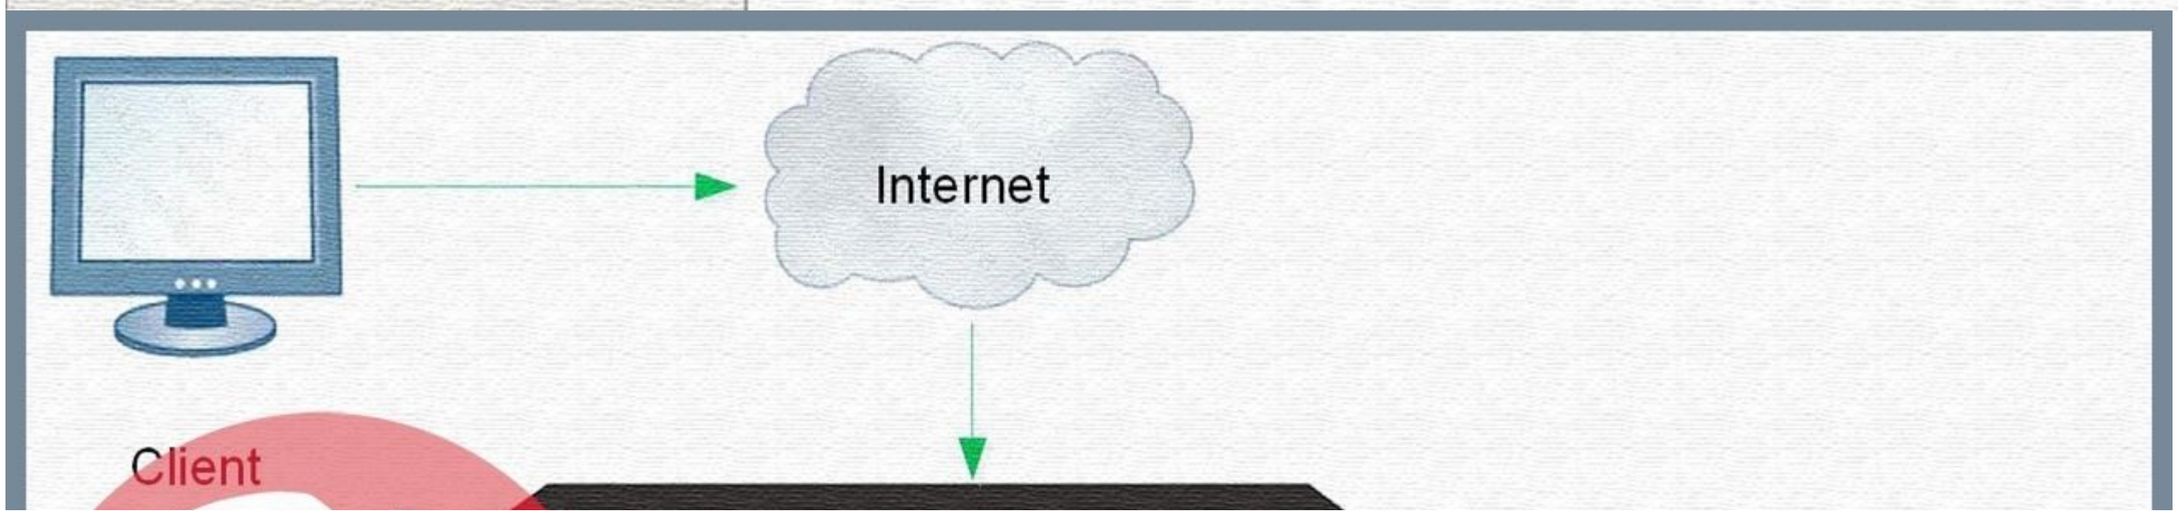

freeqas.com

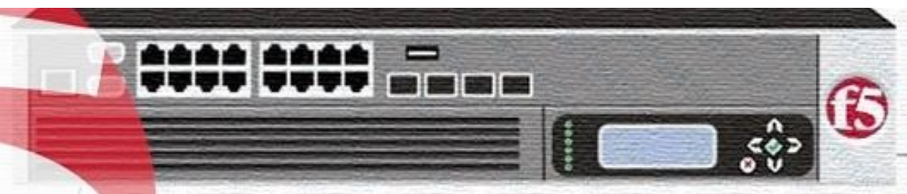

Local Traffic Manager

Ephemeral pool members:
siterequest.com:: 192.168.100.1:80
siterequest.com:: 192.168.100.2:80
siterequest.com:: 192.168.100.3:80

siterequest.com

siterequest_pool
with an FQDN pool member
*with **Auto Populate** enabled*

freeedas.com

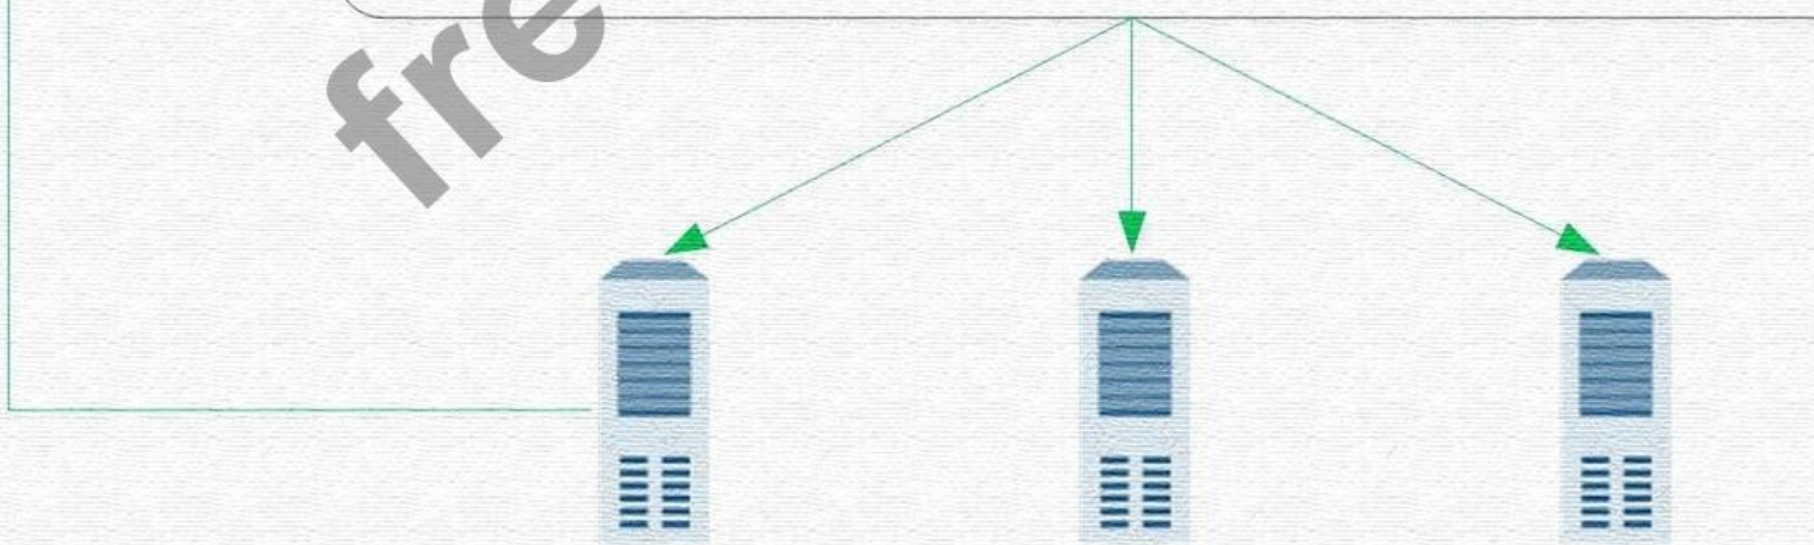

A BIG-IP pool is configured withPriority Group Activation = Less than 2 available members. The pool members have different priority groups and availability states. Which pool members arereceiving traffic?
(Choose one answer)
A. serv1
B. serv1, serv3
C. serv1, serv2, serv3, serv4
D. serv1, serv3, serv4
Answer: D (LEAVE A REPLY)
This question tests understanding ofPriority Group Activation (PGA)and how BIG-IP determines which pool members are eligible to receive traffic.
Key BIG-IP Priority Group Concepts:
* Higher priority group numbers = higher priority
* BIG-IP willonly send traffic to the highest priority groupthat meets thePriority Group Activation condition
* Lower priority groups areactivated only when the condition is met
* Onlyavailable (green)members count toward the activation threshold
Configuration from the Exhibit:
* Priority Group Activation:Less than 2 available members
* Pool Members and Status:
Pool Member
Priority Group
Status
serv1
2
Active (available)
serv2
2
Inactive (down)
serv3
1
Active (available)
serv4
1
Active (available)
Step-by-Step Traffic Decision:
* BIG-IP first evaluates thehighest priority group (Priority Group 2)
* Priority Group 2 has:
* serv1 # available

* serv2 # unavailable
* Total available members = 1
* Activation rule isLess than 2 available members
* Condition istrue(1 < 2)
* BIG-IPactivates the next lower priority group (Priority Group 1)
* Traffic is now sent to:
* serv1 (Priority Group 2)
* serv3 and serv4 (Priority Group 1)

Final Result:

Traffic is distributed toserv1, serv3, and serv4

Why the Other Options Are Incorrect:

* A- Ignores activation of the lower priority group
* B- serv4 is also active and eligible
* C- serv2 is down and cannot receive traffic

Key Data Plane Concept Reinforced:

Priority Group Activation controls when lower-priority pool members are allowed to receive traffic, based strictly on the number ofavailable membersin the higher-priority group. In this case, the failure of one high-priority member caused BIG-IP toexpand traffic distributionto lower-priority members to maintain availability.

NEW QUESTION: 19

A BIG-IP Administrator explicitly creates a traffic group on a BIG-IP device. Which two types of configuration objects can be associated with this traffic group? (Choose two.)

A. iRules
B. Application Instances
C. Virtual Addresses
D. VLANs
E. Floating Self IPs

Answer: C,E (LEAVE A REPLY)

ATraffic Groupis a collection of related configuration objects that fail over together from one BIG-IP device to another. Only "floating" objects can be members of a traffic group.

* Virtual Addresses (C):A virtual address (the IP part of a Virtual Server) is a floating object. It is assigned to a traffic group so that the entire IP moves to the standby unit during a failover.
* Floating Self IPs (E):These are used as gateways for backend servers or SNAT addresses. By associating them with a traffic group, they remain reachable by the backend network regardless of which BIG-IP is currently active.

Why other options are incorrect:

* iRules (A):iRules are configuration logic files; they are synchronized across devices but are not "hosted" by a traffic group.
* VLANs (D):VLANs are local to the hardware interfaces/trunks of each specific device and do not fail over.

NEW QUESTION: 20

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

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

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

C. Set the pool member to a Forced Offline state

D. Set the pool member to a Disabled state

Answer: C (LEAVE A REPLY)

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

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

Key behavior of pool member states:

Forced Offline

Immediately removes the pool member from load balancing.

Terminates all existing connections, regardless of persistence.

Prevents new connections from being established.

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

Disabled

Prevents new connections from being sent to the pool member.

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

Offline (non-forced)

Similar to Disabled behavior depending on context.

Does not guarantee immediate termination of existing connections.

Manually deleting connections via the command line

Is unnecessary and operationally inefficient.

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

Conclusion:

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

NEW QUESTION: 21

The diagram below shows the TCP connection setup for an application.

The diagram below shows the TCP connection setup for an application.

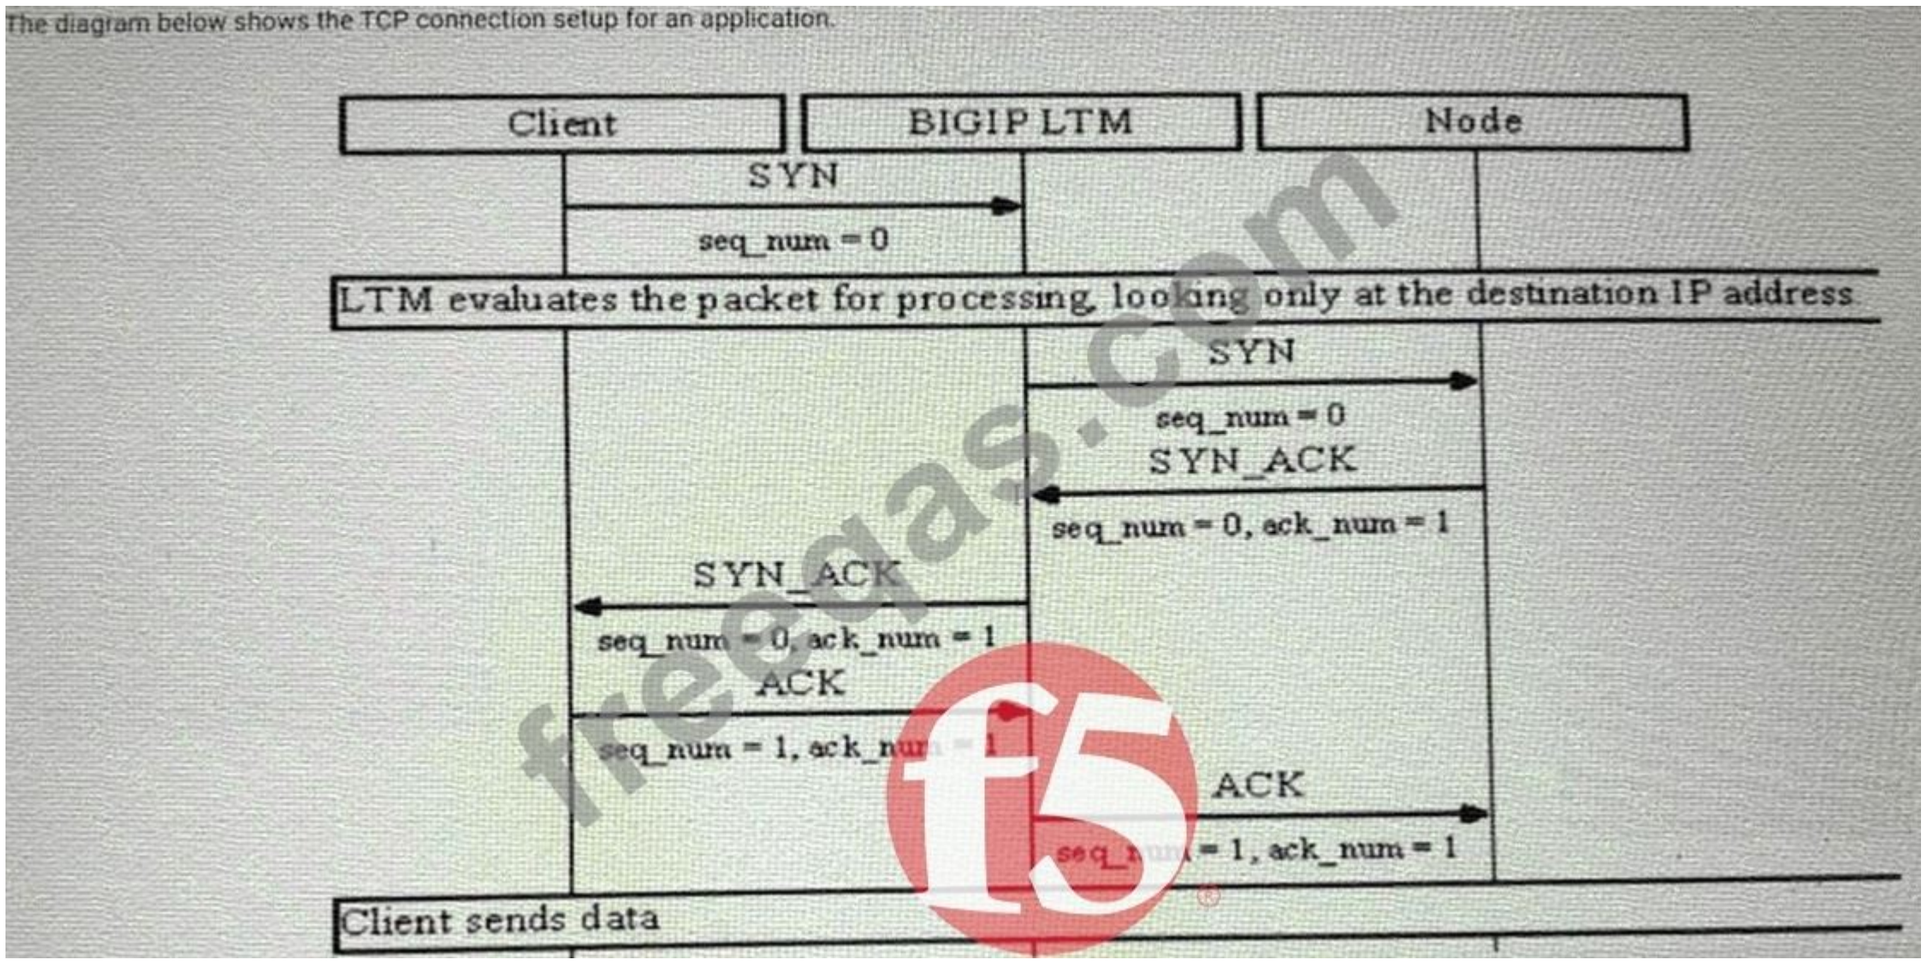

Which of the following virtual server types applies? (Choose one answer)

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 aForwarding (IP) virtual server are as follows:

* Initial Packet Processing: The diagram explicitly states that the LTM evaluates the packet lookingonly 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 aStandardvirtual server, which completes the three-way handshake with the client (SYN, SYN-ACK, ACK)beforeinitiating a separate connection to the server, the Forwarding IP virtual server passes the client's originalSYNpacket directly to the destination node.
* Response Timing: The BIG-IP system waits for theSYN-ACKfrom 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: 22
Which two statements describe differences between the active and standby systems? (Choose two.)
A. Floating selfIP addresses are hosted only by the active system.
B. Failover triggers only cause changes on the active system.
C. Virtual server addresses are hosted only by the active system.
D. Configuration changes can only be made on the active system.
E. Monitors are performed only by the active system.
Answer: A,C (LEAVE A REPLY)

NEW QUESTION: 23
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: 24
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: D (LEAVE A REPLY)
Comprehensive and Detailed Explanation From BIG-IP Administration Data Plane Concepts documents:
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: 25
The BIG-IP appliance fails to boot. The BIG-IP Administrator needs to run the End User Diagnostics (EUD) utility to collect data to send to F5 Support.
Where can the BIG-IP Administrator access this utility?
A. External VLAN interface
B. Console Port
C. Management Port
D. Internal VLAN interface
Answer: (SHOW ANSWER)

NEW QUESTION: 26
Refer to the exhibit.

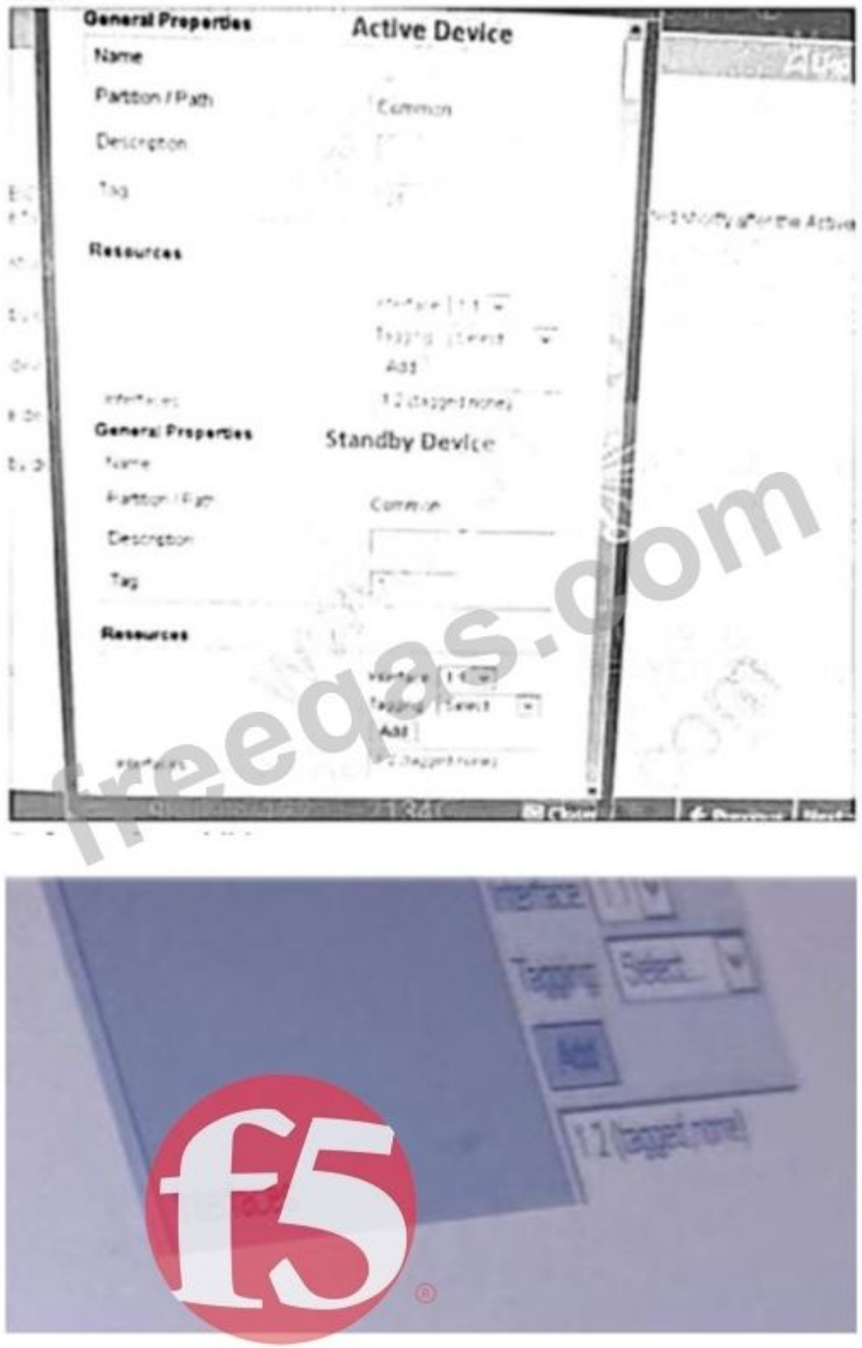

During a planned upgrade to a BIG-IP HA pair running Active/Standby, an outage to application traffic is reported shortly after the Active unit is forced to Standby. Reverting the failover resolves the outage. What should the BIG-IP Administrator modify to avoid an outage during the next failover event?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

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

freeqas.com

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

NEW QUESTION: 27
Which two statements describe differences between the active and standby systems? (Choose two.)
A. Monitors are performed only by the active system.
B. Failover triggers only cause changes on the active system.
C. Virtual server addresses are hosted only by the active system.
D. Configuration changes can only be made on the active system. (Incorrect)
E. Floating self-IP addresses are hosted only by the active system.3536
Answer: (SHOW ANSWER)
The primary distinction between Activ41e and Standby units revolves around which unit is currently processing traffic.
* Traffic Objects (C & E): The unit in the 43Active state is the only one that answers ARP requests for Virtual Server addresses and Floating Self-IPs. The Standby unit remains "quiet" for these addresses to avoid IP conflicts on the network.
* Monitors (A - False): Both the Active and Standby units perform health monitors on pool members by default. This ensures that the Standby unit is ready to take over with an up-to-date view of the pool's health.
* Failover (B - False): A failover trigger (like a VLAN fail-safe) causes the Active unit to go Standby and the Standby unit to go Active; it affects both.
* Management (D - False): Configuration changes can technically be made on either unit (though it is best practice to make them on the Active unit) and then synchronized to the peer.

NEW QUESTION: 28
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: 29
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: C (LEAVE A REPLY)
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: 30
The BIG-IP Administrator wants to provide quick failover between the F5 LTM devices that are configured in an HA Pair with a single traffic-group. The BIG-IP Administrator wants to implement the Mac Masquerade feature for this quick failover and run 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 of this, the BIG-IP Administrator must 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?
A. 02:12:34:56:01:15
B. 02:12:34:56:15:01
C. 02:12:34:56:05:dd
D. 02:12:34:56:dd:05
Answer: C (LEAVE A REPLY)
MAC Masquerade allows a traffic group to use a shared MAC address to speed up failover. When the system DB variable tm.macmasqaddr_per_vlan is enabled, the BIG-IP generates a unique MAC address for each VLAN by algorithmically modifying the base MAC address using theVLAN ID.
The calculation forVLAN 1501works as follows:
* Base MAC:The administrator set the base to 02:12:34:56:00:00.
* VLAN ID Conversion:The decimal VLAN ID (1501) must be converted into hexadecimal.
* $1501 \div 16 = 93$ remainder13(Din hex)
* $93 \div 16 = 5$ remainder13(Din hex)
* $5 \div 16 = 0$ remainder5
* Result: 1501 (Decimal) =05DD(Hex).
* Substitution:The BIG-IP replaces the last two octets of the base MAC address with the hexadecimal representation of the VLAN ID.
* Final Result:02:12:34:56:05:dd.

NEW QUESTION: 31
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: A (LEAVE A REPLY)
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.

Valid F5CAB2 Dumps shared by PrepPdf.com for Helping Passing F5CAB2 Exam! PrepPdf.com now offer the **newest F5CAB2 exam dumps**, the PrepPdf.com F5CAB2 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com F5CAB2 dumps with Test Engine here: https://www.preppdf.com/F5/F5CAB2-prepaway-exam-dumps.html (**68** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 32
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: D (LEAVE A REPLY)
Comprehensive and Detailed Explanation From BIG-IP Administration Data Plane Concepts documents:
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
Allows full use of profiles (HTTP, 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
Supports both TCP and UDP 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: 33
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: 34
Which statement is true concerning cookie persistence?
A. Cookie persistence allows persistence even if the data are encrypted from client to pool member.
B. If a client's browser accepts cookies, cookie persistence will always cause a cookie to be written to the client's file system.
C. Cookie persistence uses a cookie that stores the virtual server, pool name, and member IP address in clear text.
D. Cookie persistence allows persistence independent of IP addresses.
Answer: (SHOW ANSWER)

NEW QUESTION: 35
A BIG-IP Administrator has acluster of devices.
What should the administrator doafter creating a new Virtual Server on device 1? (Choose one answer)

A. Synchronize the settings of thegroup to device 1
B. Create a new cluster on device 1
C. Synchronize the settings ofdevice 1 to the group
D. Create a new virtual server on device 2
Answer: C (LEAVE A REPLY)
In a BIG-IPdevice service cluster, configuration objects such asvirtual servers, pools, profiles, and iRules are maintained throughconfiguration 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 alocal device
* Those changes must bepushed to the group
* The correct operation is"Sync Device to Group"
WhyCis correct:
* The virtual server was createdonly 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 mustsynchronize 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 mayremove 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 groupso all devices share the same ADC application objects.

NEW QUESTION: 36
Which three iRule events are likely to be seen in iRules designed to select a pool for load balancing? (Choose three.)
A. SERVER_SELECTED
B. SERVER_CONNECTED
C. SERVER_DATA
D. HTTP_RESPONSE
E. CLIENT_DATA
F. HTTP_REQUEST
G. CLIENT_ACCEPTED
Answer: (SHOW ANSWER)

NEW QUESTION: 37
In the GUI, where should the BIG-IP Administrator configure an existing VLAN named external to a specific interface as untagged?
A. Network - VLANs -> VLAN List -> external -> interfaces
B. Network -> VLANs -> VLAN List -> create -> external
C. Network -> VLANs -> VLAN List -> internal -> interfaces
Answer: (SHOW ANSWER)

In the BIG-IP system, VLANs are the logical entities that group physical interfaces or trunks together. To modify how a VLAN interacts with an interface, the administrator must navigate to the specific VLAN configuration object.
* VLAN List:This section displays all existing VLANs configured on the system.
* Interface Association:Within the properties of a specific VLAN (in this case, "external"), there is an Interfacessection. This is where physical ports or trunks are assigned to the VLAN.
* Tagging Status:For each associated interface, the administrator can choose betweenTagged(802.1Q) or Untagged.
* Untagged:The interface will treat incoming traffic without a VLAN header as part of this VLAN, and outgoing traffic will not have a VLAN tag added. An interface can only be "Untagged" for one VLAN.
* Tagged:Allows an interface to carry traffic for multiple VLANs simultaneously by using 802.1Q headers.

NEW QUESTION: 38
A BIG-IP Administrator is making adjustments to an iRule and needs to identify which of the 235 virtual server configured on the BIG-IP device will be affected. How should the administrator obtain this information in an effective way? (Choose one answer)
A. Local Traffic > Network Map
B. Local Traffic > iRules
C. Local Traffic > Pools
D. Local Traffic > Virtual Server
Answer: A (LEAVE A REPLY)
In a large-scale BIG-IP environment with hundreds of virtual servers, the Network Map is the most effective tool for visualizing and auditing the relationships between various ADC objects.
* The Network Map Functionality: The Network Map provides a hierarchical view of the local traffic objects. It allows an administrator to see the status and dependencies of Virtual Servers, Pools, Pool Members, and associated iRules all in one screen.
* Search and Filter: By navigating to Local Traffic > Network Map, the administrator can use the Advanced Filter. This feature allows for searching specifically for an iRule name or a string within an iRule definition. Once the filter is applied, the system displays only the Virtual Servers that are associated with that specific iRule.
* Efficiency: While the "Virtual Server List" (Option D) can be customized to show columns for iRules, it is often cumbersome to scroll through hundreds of entries. The "iRules List" (Option B) displays the scripts themselves but does not provide a reverse-lookup list of all associated virtual servers in a single view as efficiently as the Network Map.
* Summary of Relationships: The Network Map is specifically designed to answer the question, "What is this object connected to?" making it the primary administrative interface for impact analysis during configuration changes.

NEW QUESTION: 39
Refer to the exhibit.

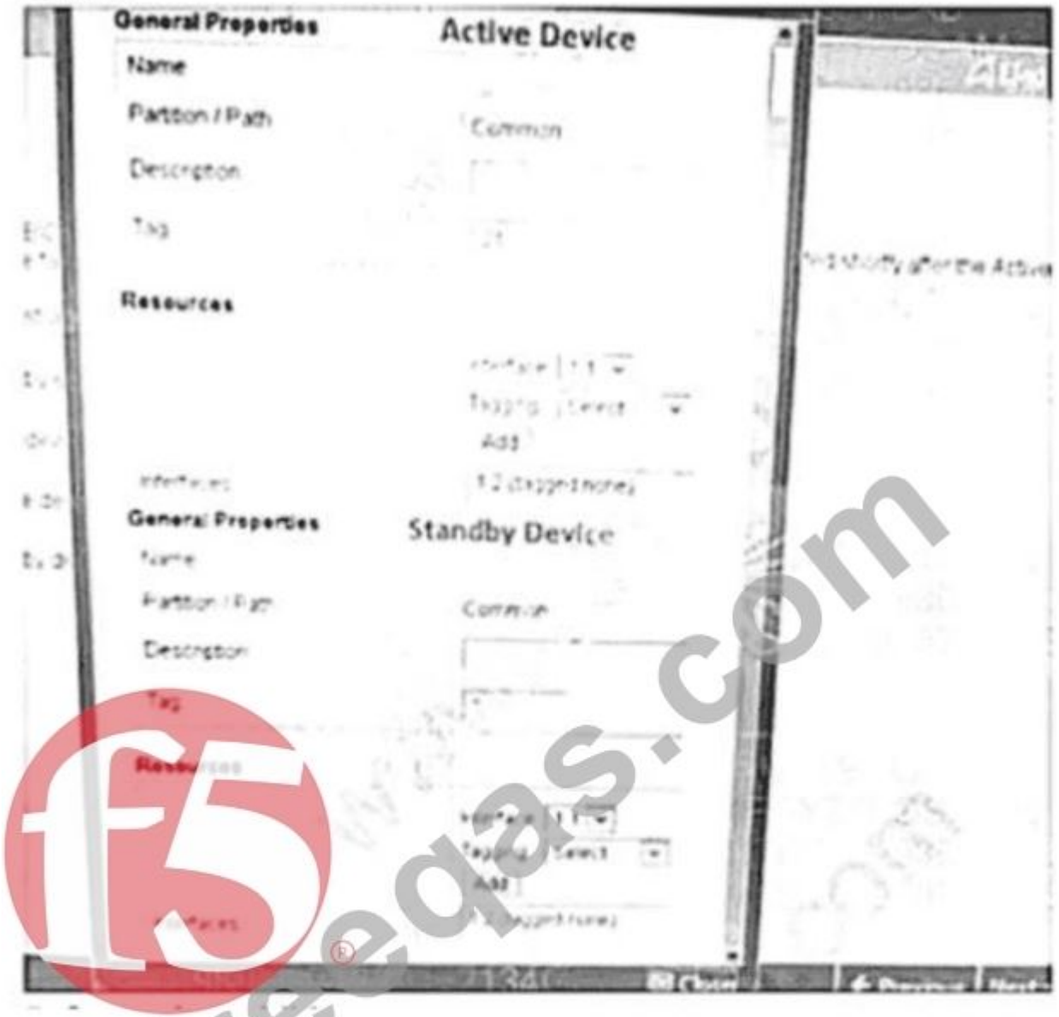

During a planned upgrade to a BIG-IP HA pair running Active/Standby, an outage to application traffic is reported shortly after the Active unit is forced to Standby. Reverting the failover

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

NEW QUESTION: 40
The diagram below shows the TCP connection setup for an application.

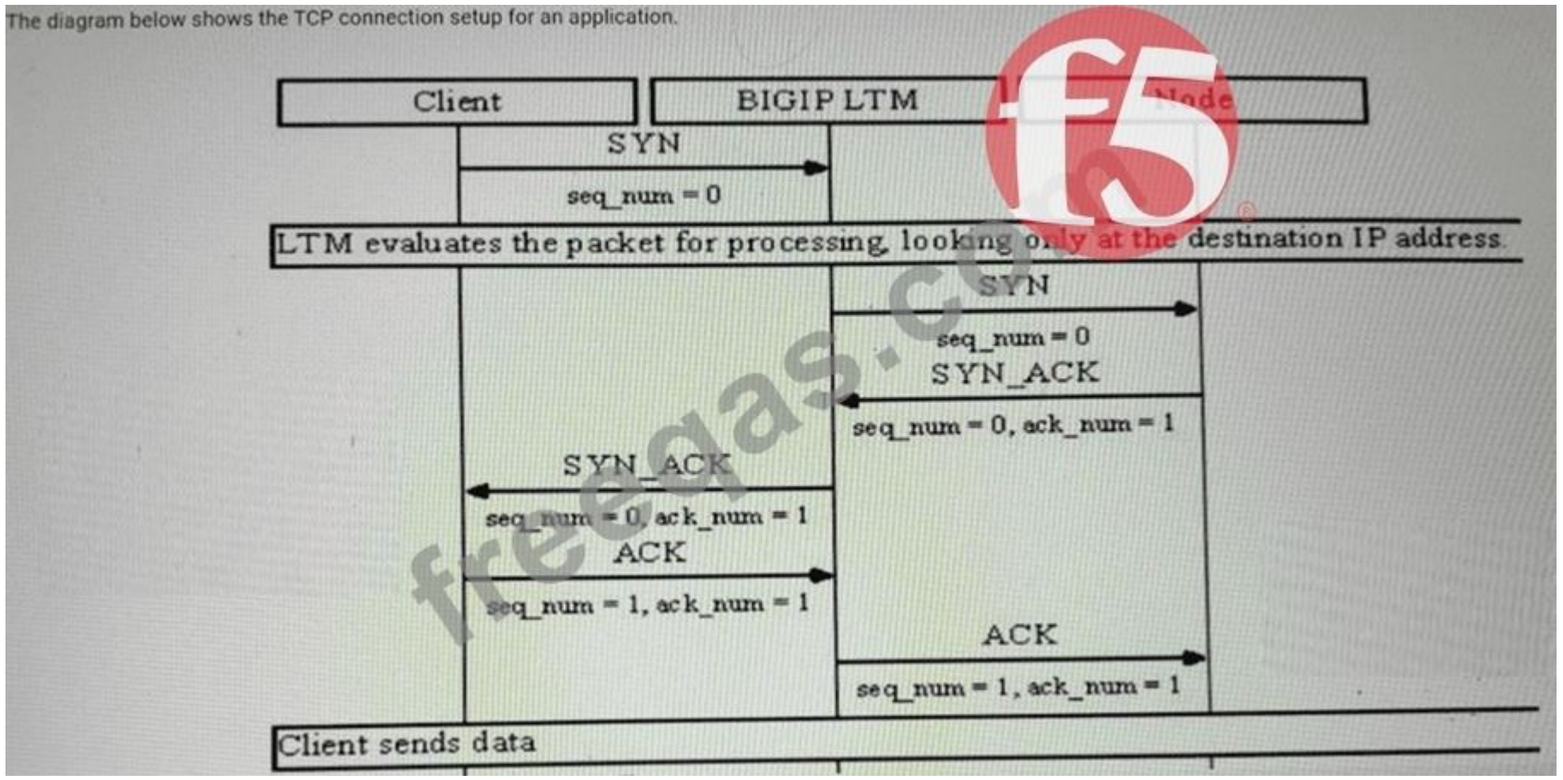

Which of the following virtual server types applies? (Choose one answer)

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: 41

Which event is always triggered when a client initially connects to a virtual server configured with an HTTP profile?

A. HTTP_REQUEST
B. CLIENT_DATA
C. HTTP_DATA
D. CLIENT_ACCEPTED
Answer: (SHOW ANSWER)

Valid F5CAB2 Dumps shared by PrepPdf.com for Helping Passing F5CAB2 Exam! PrepPdf.com now offer the **newest F5CAB2 exam dumps**, the PrepPdf.com F5CAB2 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com F5CAB2 dumps with Test Engine here: https://www.preppdf.com/F5/F5CAB2-prepaway-exam-dumps.html (**68** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)