PassQueen
a.k.a dumpsqueen.com

BIG-IP LTM Specialist: Architect Set-Up & Deploy

F5 301a

Version Demo

Total Demo Questions: 23

Total Premium Questions: 232

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Architect	10
Topic 2, Setup	49
Topic 3, Deploy	173
Total	232

QUESTION NO: 1

A virtual server is configured to handle https traffic. The clientssl profile is configured to use a2048-bit RSA key. Due to security requirements, is the LTM Specialist needs to use a

4096-bit RSA key in the future.

What two effects will this change have on the BIG-IP device? (Choose two)

A. Increase of CPU usage on the BIG-IP device

B. Decrease to 20% of licensed TPS

C. Decrease to 90% of licensed TPS

D. Increased of concurrent connection on client-side

E. Increase of TLS Renegotiation

ANSWER: A B

Explanation:

“Increase of CPU usage on the BIG-IP device” is correct because RSA private-key operations performed during a full TLS handshake are substantially more computationally expensive when the RSA modulus is increased from 2048 bits to 4096 bits. This applies particularly to client-side TLS processing, where BIG-IP must perform the server private-key operation. Even when cryptographic hardware acceleration is available, larger keys reduce the rate at which the platform can complete RSA handshakes and can increase utilization of the relevant SSL processing resources.

“Decrease to 20% of licensed TPS” is also correct. BIG-IP SSL TPS licensing normalizes capacity around 2048-bit RSA operations. A 4096-bit RSA operation consumes five times the licensed SSL TPS capacity of a 2048-bit RSA operation. Consequently, the same licensed capacity supports approximately one-fifth as many new full TLS handshakes using 4096-bit RSA keys. This is a handshake-rate capacity effect; it is not a general reduction in established connections. Capacity planning should account for the key size used by the Client SSL profile, expected new-connection/handshake rates, session reuse, and the BIG-IP platform’s SSL acceleration capabilities. F5 documents SSL/TPS licensing and cryptographic capacity considerations in its BIG-IP System Licensing documentation and provides operational TLS guidance in F5’s SSL performance guidance.

QUESTION NO: 2

Traffic to a pool of SFTP servers that share storage must be balanced by an LTM device.

What are the required profile and persistence settings for a standard virtual server?

A. tcp, ctientsst, ftp serverssl persistence

B. tcp, clientssl, serverssl persistence

C. tcp, ftp - Source address persistence

D. tcp - no persistence profile will be used

ANSWER: D

Explanation:

The required setting is **tcp - no persistence profile will be used**. SFTP is the SSH File Transfer Protocol: it operates as a subsystem over a single SSH transport connection, rather than using a separate control connection and independently load-balanced data connections. A TCP profile is therefore the applicable LTM profile for the standard virtual server. BIG-IP selects a pool member when the client TCP connection is established and maintains the connection’s flow state for its lifetime, so all packets belonging to that active SFTP/SSH session continue to reach the selected pool member without a persistence profile.

Because the SFTP servers use shared storage, a later, separate client connection can be sent to any available pool member and still access the same stored files. There is no stated application requirement to bind subsequent connections from a client to one particular server. Omitting persistence allows normal connection-by-connection load distribution while retaining the required per-connection TCP handling. SSH's single protected transport connection is defined in RFC 4253. F5's guidance on persistence profiles and their use for directing subsequent connections is available in the BIG-IP persistence profile documentation.

QUESTION NO: 3

An TLM Specialist has an Exchange that must use the LTM device to route traffic to the internet.

Which SNAT/NAT configure allows the Exchange server's traffic access the internet through the LTM device?

A. NAT

B. SNAT Pool

C. SNAT List

D. SNAT Automap

ANSWER: A

Explanation:

NAT is the appropriate configuration when the Exchange server must send traffic to the Internet through BIG-IP and requires a fixed, publicly routable translated address. A BIG-IP NAT object creates a static, one-to-one mapping between the server's original private address and a translation address. When Exchange-originated traffic is routed to BIG-IP, BIG-IP replaces the private source address with the NAT translation address before forwarding it toward the Internet. The return traffic is then translated back to the Exchange server's private address, preserving a consistent bidirectional mapping.

This is particularly appropriate for an Exchange server because externally visible services and outbound mail commonly benefit from a stable source address. The static mapping also allows the Internet-facing address to be used consistently in related DNS, mail-routing, and security policies. The translation address must be routable from the Internet-facing BIG-IP VLAN, and the upstream network must route return traffic for that address back to BIG-IP. F5 documents NAT as a static address translation mechanism, with both source and destination translation behavior. See the BIG-IP NAT tmsh reference and the BIG-IP SNAT tmsh reference.

QUESTION NO: 4

An LTM Specialist plans to enable connection mirroring for a virtual server in an HA environment.

What must the LTM Specialist consider before implementing the configuration change?

A. Impact on system performance that might be noticeable

B. The add-on license that is required for this feature to be available

C. Creating the required separate interface for connection mirroring

D. Decreased number of possible concurrent connections to that virtual server

ANSWER: A

Explanation:

Impact on system performance that might be noticeable is correct because connection mirroring replicates connection-state information from the active BIG-IP system to its peer. This state synchronization helps the standby device continue eligible flows after a failover, minimizing disruption for clients and servers. However, maintaining mirrored state requires additional processing on the traffic-management system and generates synchronization traffic between the devices. The

amount of overhead depends on factors such as the connection rate, number of connections being mirrored, traffic profile, and platform capacity.

Before enabling the feature, the LTM Specialist should therefore assess current CPU, memory, and network utilization, as well as the expected connection load during peak periods. Capacity testing or monitoring is particularly important for high-volume virtual servers, because the additional mirroring workload can become visible as reduced available performance headroom. The HA configuration must also have reliable communication between peers so that mirrored state can be delivered consistently. F5 documents connection mirroring as a high-availability feature that synchronizes connection information and notes the associated system-resource considerations. See the F5 BIG-IP LTM high-availability documentation and the F5 connection mirroring guidance.

QUESTION NO: 5

An LTM device configuration is as shown:

An LTM device configuration is as shown

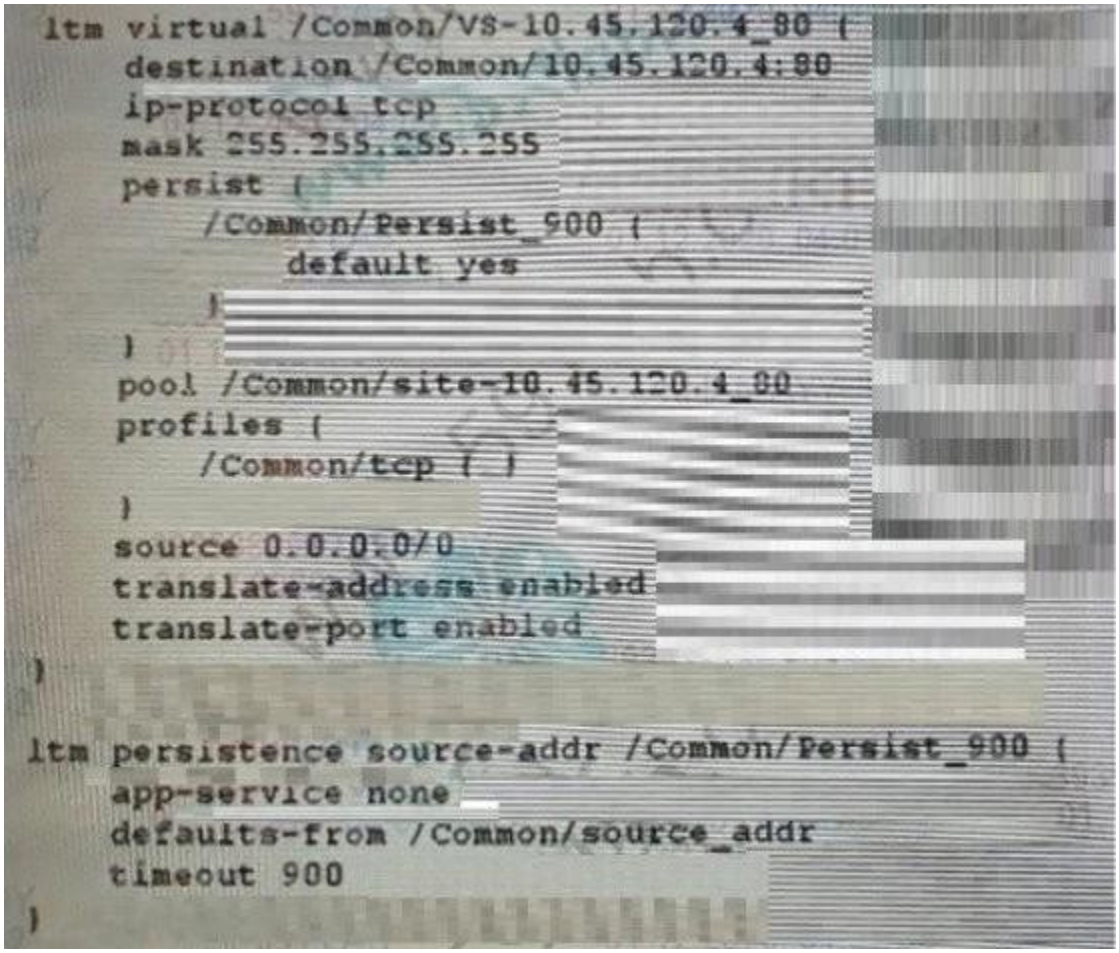

What should be the two expected outcomes based on this configuration? (Choose two.)

A. A client session that has been idle for 16 minutes will be sent to the same pool member

B. A client session that has been idle for 20 minutes will be balanced to a new pool member

C. A client session that has been idle for 14 minutes will be balanced to a new pool member

D. A client session that has been idle for 48 minutes will be sent to the same pool members

E. A client session that has been idle for 12 minutes will be sent to the same pool member

ANSWER: B E

Explanation:

The configured persistence behavior retains a client-to-pool-member association for 15 minutes after the most recent applicable connection or request. Therefore, *A client session that has been idle for 12 minutes will be sent to the same pool member* is expected because the persistence record is still valid. BIG-IP LTM uses that valid persistence record to select the previously associated pool member rather than applying the pool load-balancing method again.

A client session that has been idle for 20 minutes will be balanced to a new pool member is also expected. At that point, the configured 15-minute persistence lifetime has elapsed, so the persistence entry is no longer available to influence member selection. BIG-IP consequently treats the next connection as a new load-balancing decision and selects a pool member according to the virtual server's configured pool and load-balancing method. Persistence timeout is measured from the relevant activity that refreshes the persistence entry; continued activity can refresh the timer, while inactivity permits the record to expire. F5 documents that persistence profiles maintain client affinity for their configured timeout period and that a subsequent request is load balanced after the persistence record expires. See the F5 BIG-IP LTM persistence documentation and the F5 iRules persistence command reference.

QUESTION NO: 6

An LTM Specialist needs to provide access to a 8BG-IP to device for a company's support person access to the BIG IP device, but are NOT allowed to change any settings All support the support remote access to the BIG-IP device, but are NOT allowed to change ant settings. All support have accounts in the company's Active Directory

Which method is appropriate to provide access for the support personnel to the BIG-JP device?

A. configure remote authentication for all users with a default user role of Guest

B. configure remote authentication and map support personnel users to the Guest user role

C. configure remote authentication for all users with a default user role of Operator

ANSWER: B

Explanation:

Configure remote authentication and map support personnel users to the Guest user role. This approach lets the BIG-IP system authenticate the support staff against the company’s existing Active Directory, avoiding the need to create and maintain separate local BIG-IP accounts. Mapping only the applicable AD users or group to a BIG-IP role also applies least privilege: access is granted specifically to the support population rather than automatically assigning a role to every remotely authenticated user.

The Guest role is designed for users who need to log in and view BIG-IP configuration and status information without being permitted to make configuration changes. That meets the stated requirement that support personnel require remote access but must not be allowed to change settings. The role mapping should be implemented through the applicable remote authentication configuration, such as Active Directory or LDAP, and associated user or group role mapping. F5 documents BIG-IP user roles and their permissions in the auth user tmsh reference; remote authentication configuration is documented in the BIG-IP System Authentication Administration guide.

QUESTION NO: 7

Four members in a server pool have similar hardware platforms. An LTM Specialist needs the load balancing method that can select the server with the fewest entries in the persistence table.

Which load balancing method should the LTM Specialist use?

A. Observed

B. Dynamic Ratio

C. Least Sessions

D. Least Connections

ANSWER: C

Explanation:

Least Sessions is the appropriate method because it selects the pool member with the fewest active persistence sessions. BIG-IP maintains persistence records so that subsequent requests from a client can be directed to the same pool member when a persistence profile applies. With Least Sessions, the system uses the number of persistence-table entries associated with each eligible pool member as the selection metric, sending a new session to the member currently handling the fewest persisted sessions. This directly meets the requirement to balance according to persistence-table usage rather than according to connection count, response time, or a dynamically calculated capacity ratio. Since the pool members have similar hardware, there is no need to weight members differently; the goal is specifically to distribute persisted client sessions evenly. F5 documents Least Sessions as a dynamic load-balancing method intended for pools using persistence, and notes that it chooses the member with the fewest active sessions. See F5's BIG-IP load balancing methods reference and the BIG-IP LTM load-balancing documentation.

QUESTION NO: 8

An LTM Specialist needs to gather website statistics such as latency and throughput on the existing virtual server. This virtual server load Balances the backend web servers.

Which F5 feature will provide this?

A. the Performance panel

B. the AVR module

C. the Dashboard

D. the Statistics panel

ANSWER: B

Explanation:

The AVR module is correct because BIG-IP Analytics, formerly called Application Visibility and Reporting (AVR), is designed to collect, retain, and present application-layer traffic metrics for virtual servers. After Analytics is provisioned and an Analytics profile is associated with the virtual server, BIG-IP can measure HTTP request and response behavior and expose useful website statistics through Analytics reporting. These metrics include response time/latency-related data and throughput or traffic-volume information, allowing an LTM Specialist to assess how users experience the application while traffic is being load balanced across the pool members.

Analytics provides both real-time and historical visibility, so the specialist can review trends rather than relying only on instantaneous system counters. The collected data can also be filtered by dimensions such as virtual server, URL, client, and response status, which makes it appropriate for investigating application performance at the existing load-balanced web virtual server. F5 documents BIG-IP Analytics as the feature for monitoring application performance and traffic characteristics, including response-time and throughput reporting. See the F5 BIG-IP Analytics overview and the F5 Analytics profiles overview.

QUESTION NO: 9

An LTM Specialist needs to assign a health monitor to a pool with two pool members

10.10.10 101 and 10.10.10.102 Both pool members are listening on port 8080 with TCP. The health of the application depends on the health of an another server (10 10 10 100) that runs port 9080 with TCP.

Which two custom TCP monitors should be selected as the pool's health monitors' (Choose two)

A. a custom TCP monitor that works on port 9080 with 10.10.10.100 as alias address

B. a custom TCP monitor that works on port 9080,

C. a custom TCP monitor that works on port 8080

D. a custom TCP monitor that works on port 9080 with 10,10.10.101 as alias address

E. a custom TCP monitor that works on port 8080 with 10.10.10.101 as alias address

F. a custom TCP monitor that works on port 8080 with 10.10.10.102 is alias address

ANSWER: A C

Explanation:

The pool requires health checks for both the individual pool-member service and the shared dependency service. A custom TCP monitor on port 8080, without an alias address, uses the monitored pool member's address as its destination. When attached to this pool, BIG-IP therefore tests TCP connectivity to port 8080 on each member, 10.10.10.101 and 10.10.10.102. This confirms that the service endpoint offered by each pool member is accepting TCP connections.

A separate custom TCP monitor must use 10.10.10.100 as its alias address and 9080 as its alias port. Monitor aliases override the normal member-address destination, allowing a monitor assigned to the pool to test a fixed, external dependency instead. Because the application is considered healthy only when that dependency is reachable, this monitor ensures the dependency's TCP listener is included in the pool health decision. With multiple monitors assigned to a pool, BIG-IP evaluates the configured monitor rule; the default requirement is that all assigned monitors must report availability before a member is considered available. See the F5 TCP monitor reference and the TCP monitor API reference for alias-address and alias-port behavior.

QUESTION NO: 10

Which Standard Virtual Server settings should an LTM Specialist use to load balance across routed path of two different ISPs?

A. address translation enabled and port translation disabled

B. both address and port translation enabled

C. both address and port translation disabled

D. address translation disabled and port translation enabled

ANSWER: C

Explanation:

For a Standard Virtual Server that distributes outbound traffic across routed paths through separate ISPs, **both address and port translation disabled** is the appropriate configuration. The BIG-IP system must preserve the original destination IP address and destination service port so the packet can continue toward the actual Internet destination. The selected ISP-side path or router is used as the forwarding path, rather than becoming the translated destination endpoint.

Disabling address translation maintains transparent destination addressing. This is essential in a routed-path design because replacing the destination with a pool member address would prevent the traffic from retaining its intended external destination. Disabling port translation likewise preserves the original service port, allowing traffic for any required application or protocol to traverse the selected ISP route without a destination-port rewrite. The virtual server can therefore apply its load-balancing decision to choose the appropriate routed path while retaining the packet information needed for normal upstream routing and return traffic handling.

F5 documents address and port translation as independent virtual-server controls in the tmsh ltm virtual reference. The broader virtual-server behavior and forwarding concepts are covered in the BIG-IP Local Traffic Manager documentation.

QUESTION NO: 11

An LTM device needs an additional traffic group.

Which configuration item is required?

A. Default device

B. Group name

C. MAC Masquerade Address

D. Auto Fallback Timeout

ANSWER: B

Explanation:

Group name is required because a traffic group is a named high-availability object in BIG-IP device service clustering. The name uniquely identifies the traffic group and is the value used when associating floating objects, such as virtual addresses, with a particular active device in the device group. When an administrator creates an additional traffic group, BIG-IP must have this identifier in order to create and manage that separate ownership domain.

Traffic groups allow different floating objects to be active on different devices in a sync-failover device group, supporting active-active traffic distribution as well as controlled failover behavior. BIG-IP includes a predefined traffic group, but any additional traffic group must be created with its own name. Other traffic-group properties can be configured as needed after or during creation, depending on the desired failover design, but the named object itself cannot exist without a group name.

The BIG-IP tmsh reference defines traffic-group creation and its naming syntax in the cm traffic-group reference. F5 configuration guidance on associating traffic groups with floating objects is also available in the F5 network declarations documentation.

QUESTION NO: 12 - (SIMULATION)

A Client makes the request displayed below to the application server.

Which virtual server type should an LTM Specialist use to load balance based on the URI?

.

A.
Forwarding (Layer 2)

B.
Stateless

C.
Standard

D.
Performance (Layer 4)

ANSWER: See the explanation for the answer

Explanation:

Answer: C. Standard

A **Standard** virtual server is a full-proxy, Layer 7 virtual server type. It can use an HTTP profile and HTTP-aware policy or iRule logic to inspect the HTTP request URI and select the appropriate pool or pool member.

For example, URI-based load balancing can be implemented using an HTTP profile and an iRule such as:

```
when HTTP_REQUEST {
    if { [HTTP::uri] starts_with "/app1" } {
        pool app1_pool
```

```
    } elseif { [HTTP::uri] starts_with "/app2" } {
        pool app2_pool
    }
}
```

Forwarding (Layer 2) forwards frames and does not provide URI inspection.
Stateless does not maintain the proxy connection state needed for this Layer 7 processing.
Performance (Layer 4) is intended for fast Layer 4 traffic handling and cannot make decisions from the HTTP URI.

QUESTION NO: 13

An LTM Specialist needs to configure a setup for antivirus scanning of HTTP traffic with an internet Contact adaption Protocol (ICAP) server.

Which two server type should be used? (Choose two.)

A. Standard

B. Internal

C. Performance HTTP

D. Forwarding IP

E. Stateless

ANSWER: A B

Explanation:

Standard and **Internal** virtual servers are used in an LTM ICAP content-adaptation deployment. A Standard virtual server provides the full-proxy HTTP processing required to inspect client requests, apply an HTTP profile, and invoke a request-adaptation or response-adaptation profile. The adaptation profile sends the applicable HTTP content to the configured ICAP service, allowing the antivirus engine to evaluate the content before BIG-IP continues processing the transaction. This full-proxy behavior is essential because BIG-IP must be able to parse and manage the HTTP flow while coordinating the adaptation exchange.

An Internal virtual server supports the local traffic-processing portion of the design. Internal virtual servers accept traffic that originates within BIG-IP rather than traffic arriving directly on a VLAN. This makes them appropriate where the content-adaptation workflow or related local service chaining requires traffic to be handled internally by the BIG-IP system. Together, these virtual-server types provide the client-facing full-proxy HTTP entry point and the internal processing path needed for the ICAP antivirus integration. F5's content-adaptation guidance describes associating ICAP request or response adaptation with HTTP traffic processing: F5 BIG-IP Content Adaptation documentation. F5 also documents the role and behavior of virtual servers, including internal virtual servers, in its BIG-IP documentation.

QUESTION NO: 14

A company plans to launch a huge marketing campaign and expects increase demand of their secure website. With the current virtual server setup, the LTM Specialist expects that the LTM device will reach its capacity limits. For the wen application to function properly. Cookies persistence is required. The LTM Specialist needs to reduce LTM device load without breaking the application.

Which two settings should the LTM Specialist modify to meet the requirement? (Choose two.)

A. Remove HTTP compression profile

B. Remove HTTP profile

C. Remove web acceleration profile.

D. Modify virtual Server type to performance (Layer 4)

E. Remove ClientSSL profile

ANSWER: A C

Explanation:

Removing the HTTP compression profile and removing the web acceleration profile are appropriate ways to reduce BIG-IP processing load while preserving the application requirements. HTTP compression requires the system to inspect eligible HTTP responses and perform compression work before sending them to clients. Disabling it removes that per-response processing overhead, at the cost of potentially using more client-side bandwidth. F5 documents the HTTP Compression profile as the mechanism used to compress HTTP server responses, so it is an optional optimization rather than a requirement for cookie-based persistence.

Similarly, the Web Acceleration profile enables application-delivery optimizations such as caching and content transformations. These features can consume system resources, particularly during a large traffic increase. Removing this profile reduces the processing associated with those optional acceleration services without eliminating the HTTP processing needed for cookie persistence. Cookie persistence depends on HTTP-aware handling so that BIG-IP can inspect the relevant HTTP cookie data; therefore, keeping the HTTP profile supports the stated application requirement. The secure-site design can also retain client-side SSL processing and a standard virtual server while these optional performance features are removed. See F5 documentation for HTTP Compression profiles and Web Acceleration profiles.

QUESTION NO: 15

An application is sensitive to packet loss and unexpected session termination. A pair of LTM devices is configured in an Active/Standby high availability configuration. SNATS are

NOT used and the virtual server contains a Universal Persistence profile.

which two actions must an LTM Specialist take to ensure the sessions are maintained between the client and server during an LTM device failover event while maintaining maximum uptime? (Choose two.)

A. configure a serial failover cable for mirror traffic

B. configure a OneConnect profile to mirror connections

C. configure a VLAN and primary mirroring address for mirror traffic

D. enable Mirroring for a virtual server and persistence profile

E. enable Clone Pools for a virtual server and a persistence profile

ANSWER: C D

Explanation:

To preserve existing application sessions through an Active/Standby failover, BIG-IP must synchronize both its connection state and the Universal Persistence records to the peer device. Enabling mirroring on the virtual server causes the active device to mirror eligible connection state to the standby device. When the standby takes over, it can continue processing the established client-side and server-side flows rather than treating them as new connections. This is especially important for an application that cannot tolerate unexpected session termination.

Persistence mirroring must also be enabled on the Universal Persistence profile. Universal Persistence commonly stores application-specific session mappings, and mirroring those records ensures that the newly active device has the persistence state required to select the same pool member for subsequent requests. Together, connection mirroring and persistence mirroring provide continuity for in-flight connections and for ongoing persistence decisions.

A dedicated mirroring VLAN with a primary mirroring address provides the peer-to-peer network path used to carry mirrored state. Each device requires the appropriate mirroring configuration so the active unit can send state updates to its peer. F5 documents the connection and persistence mirroring requirements in K13310: Overview of connection and persistence mirroring and the configuration concepts in the BIG-IP Connection Mirroring documentation.

QUESTION NO: 16

An LTM Specialist manages an Active/Standby high availability pair. A device must fail over when all interfaces associated with a particular external VLAN become unavailable, even if communication with the peer device remains available.

Which configuration item should the LTM Specialist enable?

A. VLAN Fail-safe

B. Gateway Fail-safe

C. Connection Mirroring

D. Network Failover

ANSWER: A

Explanation:

VLAN Fail-safe is correct because it monitors the availability of interfaces associated with a VLAN. When the configured VLAN is considered unavailable, BIG-IP can take the configured fail-safe action, including failing over the active traffic group to the peer device. This protects against a condition where the active device remains reachable by its peer but has lost connectivity to the client-side or network-side VLAN required to process application traffic.

VLAN fail-safe is configured on the affected VLAN and is separate from gateway fail-safe. Gateway fail-safe tests reachability of a configured gateway, whereas VLAN fail-safe reacts to loss of the VLAN's interfaces. F5 documents VLAN fail-safe behavior in the BIG-IP Device Service Clustering Administration documentation.

QUESTION NO: 17

Exhibit.

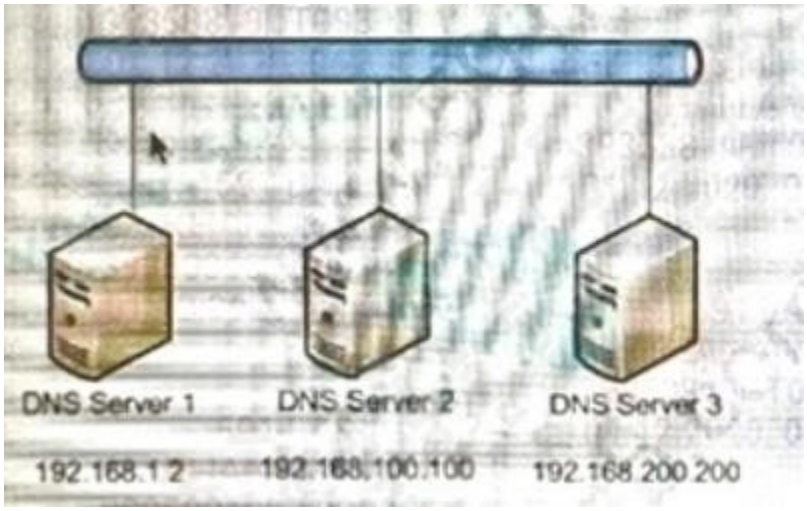

The server team has recently configured the three new DNS servers shown for the data center. No current DNS servers are currently configured on the LTM device.

Which command should be used to configure the LTM device to use the new DNS servers?

A. tmsh create/systins name-servers add {192
168.1.2.192.168.100.1O0.192.168.2O0.2OO}

B. tmsh change /sysdns name-servers add
{192.168.1.2.192.168.100.100.192.168.200.200}

C. tmsh modify /sys dns name-servers add { 192.168.1.2 192.168.100.100 192.168.200.200 }

D. tmsh set/sysdns name-servers add {192.168.2.192.168.100.100.192.168.200.200}

ANSWER: C

Explanation:

The command `tmsh modify /sys dns name-servers add { 192.168.1.2 192.168.100.100 192.168.200.200 }` is correct because DNS resolver settings on BIG-IP are configured through the system DNS object, addressed in TMSH as `/sys dns`. This is a system-level configuration object that exists independently of whether name servers have already been configured. Therefore, adding resolver addresses to an initially empty `name-servers` property is performed with `modify`, not by creating a separate DNS object.

The `name-servers add` syntax appends the supplied IPv4 addresses to the BIG-IP system resolver configuration. The addresses must be individual, space-separated entries within braces. Once configured, the BIG-IP system can use these resolvers for system DNS lookups, including lookups needed by features that rely on the host resolver configuration. F5 documents the `sys dns` TMSH component and its `name-servers` property in the TMSH reference: F5 TMSH sys dns reference. Additional system DNS configuration guidance is available in the BIG-IP DNS settings documentation.

QUESTION NO: 18

Two client networks use overlapping addresses. Clients from each network must access a virtual server using the destination address 10.50.10.25, but traffic from each network must remain isolated.

Which two actions should the LTM Specialist take? (Choose two.)

A. Assign each client-facing VLAN to a different route domain

B. Create each virtual server destination with its route domain identifier

C. Configure SNAT Automap on both virtual servers

D. Configure both client VLANs in the default route domain

E. Enable address translation on the virtual servers

ANSWER: A B

Explanation:

The LTM Specialist must assign the client-facing VLANs to **different route domains** and create the virtual server destinations using the appropriate **route domain identifiers**. Route domains permit BIG-IP to maintain separate routing and traffic-processing tables for overlapping IP address spaces. A destination such as `10.50.10.25%1` is distinct from `10.50.10.25%2`, even though the base IPv4 address is identical.

VLAN separation alone does not provide complete route isolation for overlapping addresses. SNAT can change source addresses but does not define independent routing tables for identical client or destination networks. F5 documents route-domain behavior and route-domain identifiers in the BIG-IP Network Administration documentation.

QUESTION NO: 19

A customer wants to select the pool for an application based on information found in the path of the URL.

For example:

http://www.example.com/app 1 should be sent to the app 1 pool

http.//www.exampie.com/app 2 should be sent to the app2 pool

Which two profiles need to be assigned to the virtual server? (Choose two.)

A. Client SSL

B. Persistence

C. TTP Compression

D. HTTP

E. TCP

ANSWER: D E

Explanation:

The **HTTP** and **TCP** profiles are required for this design. The HTTP profile enables BIG-IP LTM to parse client HTTP requests, including the request URI and its path component. Once HTTP processing is enabled, a traffic policy or iRule can evaluate a URI such as `/app1` or `/app2` and select the corresponding pool. For example, an iRule can use the `HTTP_REQUEST` event and the `HTTP::uri` command to inspect the requested URI, then direct the request to the appropriate application pool.

The TCP profile supplies the full-proxy TCP connection handling needed by a standard HTTP virtual server. BIG-IP accepts the client-side TCP connection, processes the HTTP request at Layer 7, makes the pool-selection decision based on the path, and establishes or uses a server-side TCP connection to the selected pool member. This combination supports path-based application routing while retaining standard TCP transport behavior. F5 documents URI access through HTTP::uri and describes HTTP request processing in its HTTP_REQUEST event reference.

QUESTION NO: 20

Exhibit.

The LTM devices LTM1 and LTM2 are configured in a Device Group (Sync Failover) with Network Failover configured on both the management and HA and Internal VLANS. and ConfigSync is confined in a Device Group (Sync Failover) with Network Failover and internal are tagged on a single trunk with subnets Connection Mirroring is configured on both the HA interlace directly connected between LTM1 and LTM2, and the management interlace is connected to a management switch. The LTM devices have four Traffic Groups defined, and both LTM devices are healthy and capable of passing traffic for any of the Traffic Groups.

An LTM Specialist disconnects the cable for the HA network in an effort to test failover.

Which HA functionality works in this case?

A. ConfigSync does NOT work. Connection Mirroring does NOT work.

B. ConfigSync works. Connection Mirroring works.

C. ConfigSync works. Connection Mirroring does NOT work.

D. ConfigSync does NOT work; Connection Mirroring works.

ANSWER: C

Explanation:

ConfigSync works. Connection Mirroring does NOT work is correct because the directly connected HA network cable is the configured transport path for connection mirroring. Connection mirroring continuously sends connection-state updates between the devices across the configured mirror address and port. Disconnecting that physical HA link removes the path used to replicate those state records, so newly established or updated connection state cannot be mirrored across that interface.

ConfigSync is a separate device-service function from connection mirroring. In this topology, the management interfaces remain connected through the management switch, providing a remaining management-network path between the peers. Therefore, the devices can still exchange configuration synchronization traffic through their available ConfigSync connectivity even though the dedicated HA cable is unavailable. Traffic-group ownership and failover communication are also distinct from the connection-mirroring data channel; this question specifically distinguishes the surviving ConfigSync path from the failed mirroring path.

F5 documents connection mirroring as a mechanism that mirrors connection state between peer devices and requires network connectivity between the configured mirror endpoints. See F5 BIG-IP Device Service Clustering: Configuring connection mirroring and F5: ConfigSync and device service clustering overview.

QUESTION NO: 21

Two BIG-IP devices are configured in a sync-failover device group. The LTM Specialist needs one application virtual address to be active on the first device and another application virtual address to be active on the second device. Each application must fail over to the peer device if its current active device becomes unavailable.

Which two actions should the LTM Specialist take? (Choose two.)

A. Create an additional traffic group

B. Assign each virtual address to the appropriate traffic group

C. Create separate device groups for the two virtual servers

D. Enable connection mirroring on both virtual servers

E. Configure different SNAT pools for each application

ANSWER: A B

Explanation:

The LTM Specialist must **create an additional traffic group**. Traffic groups define independent failover ownership domains for floating objects. A sync-failover device group can therefore support active-active operation by allowing separate traffic groups to be active on different devices.

The LTM Specialist must also **assign each virtual address to the appropriate traffic group**. Virtual addresses are floating objects, and their traffic-group association determines which BIG-IP device owns and advertises each address. When a device fails, the traffic group and its associated virtual address move to the peer.

Assigning different virtual servers to separate pools does not distribute floating ownership. The traffic-group association is made on the virtual address. F5 documents traffic groups in the tmsh traffic-group reference.

QUESTION NO: 22

AN LTM Specialist needs to determine the delay between an LTM device and the internal web server for a specific client.

Which two AVR reporting options should the LTM Specialist enable to measure the delay?

(Choose two.)

A. User agents

B. Methods

C. Response codes

D. Server latency

E. Client IP

ANSWER: D E

Explanation:

Server latency and **Client IP** must be enabled in the HTTP Analytics profile. Server latency is the relevant timing metric because it measures the time associated with the server-side portion of the transaction: BIG-IP sends the request toward the

pool member and receives the response back. This allows the specialist to evaluate the delay between the LTM device and the internal web server, rather than client-network or complete page-load time.

Enabling Client IP makes the client address available as an Analytics reporting dimension. The specialist can then use that dimension to isolate or filter the reported server-latency data for the particular client in question. Together, these settings provide both the required measurement and the client-specific context needed for the investigation. BIG-IP Analytics reporting is driven by the metrics and dimensions explicitly selected in the Analytics profile, so both must be configured before the desired client-level server-delay report can be produced. F5 describes Analytics profiles and configurable collected statistics in its BIG-IP Analytics implementation documentation. Additional reporting and profile configuration guidance is available in the AVR setup documentation.

QUESTION NO: 23

An LTM Administrator receives an email from the NOC stating that the switch connected to the backend server was shut down for maintenance. The BIG-IP device handles only UDP traffic. The BIG IP device did not fail over to a DR location when no pool members were available.

When the LTM Administrator checks the pool, it confirms that the monitor is still marking UP the pool member.

A tcpdump of the traffic shows the following output:

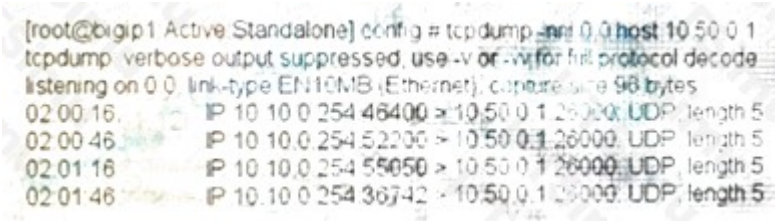

A list of the monitor configuration shows the following:

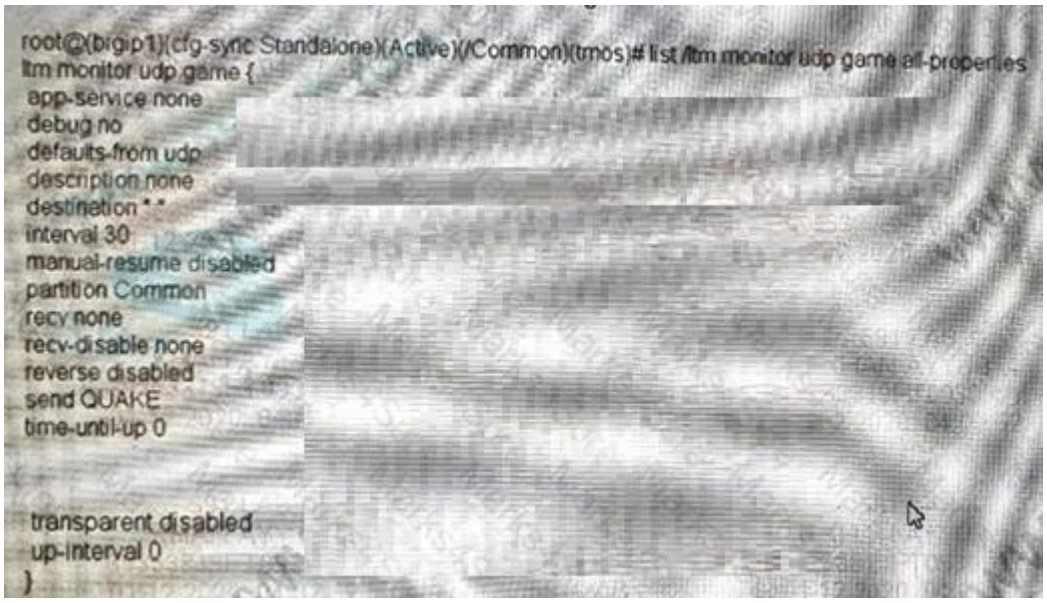

Which two modifications to the LTM configuration will mark this pool member down, when the switch is down? (Choose two.)

A. increase the timeout to three times the interval

B. add a receive string to the game monitor

C. enable reverse and wait for the next connection

D. also assign a gatewayjcmp monitor to the pool

E. enable manuai-resume on the same monitor

ANSWER: B D

Explanation:

Adding a receive string to the game monitor makes the UDP health check depend on an affirmative application response. A UDP monitor without a receive string can consider a target available when it does not receive an ICMP error, because UDP itself has no connection establishment or acknowledgement mechanism. When the access switch is unavailable, traffic can simply be dropped; in that case, the lack of an ICMP response alone does not prove that the UDP service is healthy. Configuring a receive string causes the monitor to require the expected response within its timeout, so the pool member is marked down when the switch outage prevents the server response from returning.

Assigning a gateway ICMP monitor provides an additional network-path check. The gateway ICMP monitor verifies reachability of the next-hop gateway used to reach the pool member. If the relevant gateway or path is unreachable during the outage, the monitor fails and contributes to the member being marked unavailable. This is particularly useful for distinguishing application-level UDP behavior from underlying network reachability. F5 documents the UDP monitor receive-string behavior in the UDP monitor tmsh reference and gateway reachability monitoring in the gateway ICMP monitor tmsh reference.