DUMPSARENA

LTM Specialist: Maintain & Troubleshoot

F5 301b

Version Demo

Total Demo Questions: 32

Total Premium Questions: 328

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Troubleshooting	240
Topic 2, Maintenance	88
Total	328

QUESTION NO: 1

An LTM HTTP pool has an associated monitor that sends a string equal to 'GET /test.html'.

Which two configurations could an LTM Specialist implement to allow server administrators to disable their pool member servers without logging into the LTM device? (Choose two.)

A. Set monitor to transparent and ask the server team to set string ‘TRANSPARENT’ in test.html.

B. Set ‘receive string’ equal to 'SERVER UP and ask the server team to set string ‘SERVER DOWN’ in test.html.

C. Set ‘alias’ equal to 'SERVER DOWN’ and ask the server team to set string ‘SERVER DOWN’ in test.html.

D. Set ‘receive disable string’ equal to 'SERVER DOWN’ and ask the server team to set string ‘SERVER DOWN’ in test.html.

E. Set ‘disable pool member’ equal to 'SERVER UP’ and ask the server team to set string ‘SERVER DOWN’ in test.html.

ANSWER: B D

Explanation:

Setting the receive string to `SERVER UP` and having administrators change the content of `test.html` to `SERVER DOWN` provides an application-controlled withdrawal mechanism. An HTTP monitor considers its probe successful only when the expected receive string is found in the HTTP response. Once the expected `SERVER UP` text is no longer returned, the monitor fails and BIG-IP marks the pool member unavailable, preventing new connections from being selected for it. This lets the server team control availability through the served health-check content.

Using a receive disable string of `SERVER DOWN` is the explicit monitor feature for graceful administrative removal. When that string is present in the HTTP response, the monitor identifies the member as disabled rather than simply unavailable. BIG-IP then avoids sending new traffic to that member while retaining the meaningful disabled state, which is useful when administrators intentionally place an application instance into maintenance. The health-check endpoint can therefore act as an operational switch managed by the server administrators, with no BIG-IP login required. F5 documents the HTTP monitor `recv` and `recv-disable` properties in its tmsh HTTP monitor reference; related monitor behavior is described in the BIG-IP health monitor documentation.

QUESTION NO: 2

-- Exhibit –

-- Exhibit --

Refer to the exhibits.

An LTM Specialist has configured a virtual server to distribute connections to a pool of application servers and to offload SSL processing. The application fails to work as expected when connecting to the virtual server. It does work when clients connect directly to the application. Two packet captures were taken at the application server.

What is the root cause of the problem?

A. The application servers are NOT listening on port 80.

B. The LTM device is sending non-SSL traffic to an SSL port.

C. The virtual server does NOT have a clientSSL profile assigned.

D. The SSL handshake between the LTM device and the server is failing.

ANSWER: B

Explanation:

The LTM device is sending non-SSL traffic to an SSL port. With SSL offload, a Client SSL profile terminates the client-side TLS connection on BIG-IP. After decryption, BIG-IP forwards the resulting cleartext application traffic—typically HTTP—to

the pool member unless a Server SSL profile is also configured to establish TLS for the server-side connection. The packet captures at the application server therefore reveal the protocol mismatch: the server is receiving unencrypted HTTP on a destination port that expects an SSL/TLS handshake, commonly TCP 443. Because the application works when accessed directly, the application server and its TLS service are functional; the failure occurs specifically in the BIG-IP-to-server leg after client-side SSL termination. The appropriate design is either to direct the pool to the server's cleartext HTTP listener, such as port 80, or to configure server-side SSL re-encryption with a Server SSL profile when the pool members must continue receiving HTTPS. F5 documents that Client SSL profiles handle client-side SSL/TLS processing, while Server SSL profiles provide SSL/TLS processing for connections from BIG-IP to pool members. See the F5 SSL Profiles Reference and F5 SSL iRules documentation.

QUESTION NO: 3

A new web application is hosted at www.example.net, but some clients are still pointing to the legacy web application at www.example.com.

Which iRule will allow clients referencing www.example.com to access the new application?

A. when HTTP_REQUEST {

B. when HTTP_REQUEST { if { [HTTP::host] eq "www.example.com" } { HTTP::host "www.example.net" } }

C. when HTTP_DATA {

D. when HTTP_RESPONSE {

ANSWER: B

Explanation:

The correct iRule is `when HTTP_REQUEST { if { [HTTP::host] eq "www.example.com" } { HTTP::host "www.example.net" } }`. The `HTTP_REQUEST` event runs after BIG-IP receives the client's HTTP request and before it forwards that request to the selected pool member. This provides the appropriate point to inspect the incoming `Host` header and replace the legacy hostname with the hostname expected by the new application. The client can continue sending requests for `www.example.com`, while the backend application receives `Host: www.example.net` and can therefore route the request to the correct virtual host or application configuration. The URI, method, request body, and other request headers are left intact, so this approach transparently preserves the client's requested resource path. F5 documents that `HTTP::host` returns or sets the HTTP Host header, and that setting it within an iRule changes the header passed upstream. The HTTP request-processing event is documented by F5 as the request-stage event for HTTP traffic. See the F5 iRules references for HTTP::host and HTTP_REQUEST.

QUESTION NO: 4

The pool members are serving up simple static web content.

The current virtual server configuration is given as follows:

tmsh list ltm virtual simple

ltm virtual simple {

destination 10.10.10.10:80

ip-protocol tcp

mask 255.255.255.255

profiles {

http { }

httpcompression { }

oneconnect { }

tcp { }

}

snat automap

vlans-disabled

}

tmsh list ltm pool simple_pool

ltm pool simple_pool {

members {

10.10.10.11:80 {

address 10.10.10.11 }

10.10.10.12:80 {

address 10.10.10.12 }

10.10.10.12:80 {

address 10.10.10.13 }

}

}

Which three objects in the virtual server configuration can be removed without disrupting functionality of the virtual server? (Choose three.)

A. tcp

B. http

C. oneconnect

D. snat automap

E. httpcompression

ANSWER: B C E

Explanation:

The **http**, **oneconnect**, and **httpcompression** profiles can be removed together for this virtual server because none is required to load balance basic TCP connections to pool members serving static content. The HTTP profile enables BIG-IP to parse HTTP messages and apply HTTP-layer features. If the virtual server does not need HTTP request inspection, header manipulation, HTTP redirects, or other Layer 7 processing, traffic can be passed through as TCP without that profile.

The OneConnect profile is an optimization that permits server-side connection reuse across compatible client requests. It is useful for improving connection efficiency in certain HTTP workloads, but it is not a prerequisite for delivering requests to pool members.

The HTTP compression profile is likewise an optional content transformation feature. It relies on HTTP processing to identify suitable responses and apply compression. Removing the HTTP and HTTP compression profiles as a pair simply means BIG-IP no longer parses or compresses HTTP payloads; the origin servers can still return their static responses normally. This leaves the virtual server able to establish and proxy the TCP flows required for the web service. F5 describes the purposes and configuration of these LTM profiles in the LTM Profiles Reference and explains virtual-server traffic processing in the BIG-IP virtual server documentation.

QUESTION NO: 5

Users in a branch office are reporting a website is always slow. No other users are experiencing the problem. The LTM Specialist tests the website from the external VLAN along with testing the servers directly. All tests indicate normal behavior. The environment is a single HTTP virtual server on the external VLAN with a single pool containing three HTTP pool members on the internal VLAN.

Which two locations are most appropriate to collect additional protocol analyzer data? (Choose two.)

A. a user's machine

B. the switch local to the user

C. the LTM device's internal VLAN

D. the LTM device's external VLAN

E. a user's Active Directory authentication

ANSWER: A B

Explanation:

Capturing traffic at **a user's machine** and **the switch local to the user** is appropriate because the reported behavior is limited to one branch office, while tests initiated at the BIG-IP external VLAN and directly against the servers are normal. These capture points place the investigation at the affected client edge and on the branch network segment immediately adjacent to that client population. A client-side trace can show browser-to-virtual-server timing, DNS resolution behavior, TCP retransmissions, receive-window conditions, and application response timing as experienced by the user. A simultaneous capture or SPAN on the local switch helps determine whether delay, loss, retransmission, duplex issues, congestion, or asymmetric handling occurs before traffic reaches the BIG-IP system. Comparing packet timestamps and TCP sequence behavior between these two locations is a practical way to localize the problem to the endpoint, its local access network, or the upstream path. F5 documents packet capture as a core troubleshooting technique and supports capturing traffic with `tcpdump` when analysis must continue at the BIG-IP boundary. See F5: Overview of packet tracing with tcpdump and F5 BIG-IP network traffic troubleshooting.

QUESTION NO: 6

An LTM Specialist wants an HTTP virtual server to maintain session persistence by inserting a BIG-IP cookie into server responses.

Which two profiles must be associated with the virtual server? (Choose two.)

A. An HTTP profile

B. A cookie persistence profile configured for cookie insert

C. A OneConnect profile

D. An HTTP compression profile

E. A Server SSL profile

ANSWER: A B

Explanation:

An HTTP profile and **a cookie persistence profile configured for cookie insert** are required. The HTTP profile enables BIG-IP to parse HTTP requests and responses, which is necessary before it can inspect and insert HTTP headers. The cookie persistence profile configured for the cookie insert method instructs BIG-IP to add its persistence cookie to an eligible HTTP response and use that cookie on subsequent requests.

A TCP profile is normally present for TCP virtual servers, but it does not provide HTTP message parsing or cookie insertion. A OneConnect profile can improve server-side connection reuse for HTTP traffic, but it is not required for cookie persistence. F5 documents cookie persistence configuration in the tmsh cookie persistence profile reference.

QUESTION NO: 7

A virtual server for a set of web services is constructed on an LTM device. The LTM Specialist has created an iRule and applied this iRule to the virtual server:

when HTTP_REQUEST {

switch [HTTP::uri] {

"/ws1/ws.jsp" {

log local0. "[HTTP::uri]-Redirected to JSP Pool"

pool JSP

}

default { log local0. "[HTTP::uri]-Redirected to Non-JSP Pool"

pool NonJSP

}

}

}

However, the iRule is NOT behaving as expected. Below is a snapshot of the log:

/WS1/ws.jsp-Redirected to JSP Pool

/WS1/ws.jsp-Redirected to JSP Pool

/WS1/ws.jsp-Redirected to JSP Pool

/WS1/WS.jsp-Redirected to Non-JSP Pool

/ws1/WS.jsp-Redirected to Non-JSP Pool

/WS1/ws.jsp-Redirected to JSP Pool

/ws1/ws.jsp-Redirected to Non-JSP Pool

What should the LTM Specialist do to resolve this?

A. Use the following: `switch -lc [HTTP::uri]`

B. Use the following: `switch [string tolower [HTTP::uri]]`

C. Set the "Case Sensitivity" option of each member to "None".

D. Select the "Process Case-Insensitivity" option for the virtual server.

ANSWER: B

Explanation:

Use `switch [string tolower [HTTP::uri]]`. The `HTTP::uri` command returns the request URI as received by BIG-IP, including the path and any query string, and URI path matching in a Tcl `switch` is case-sensitive unless the value or matching behavior is explicitly normalized. The log entries show the intended JSP resource being requested with different capitalization, such as `/WS1/ws.jsp` and `/ws1/WS.jsp`. Converting the URI to lowercase before evaluating it creates one

consistent value for all capitalization variants. Consequently, each of those path variants becomes `/ws1/ws.jsp`, which matches the iRule's JSP case and sends the request to the JSP pool. This change should be made at the expression being tested by `switch`, while the existing URI may still be logged if retaining the original client-supplied casing is useful for troubleshooting. F5 documents that `HTTP::uri` supplies the HTTP request URI for iRule processing, and Tcl documents that `string tolower` converts a string to lowercase. See F5 HTTP::uri documentation and the Tcl string command reference.

QUESTION NO: 8

An LTM Specialist is troubleshooting an issue where one LTM device in a three LTM device group is failing to synchronize after a synchronize to group command is issued. The LTM Specialist verifies there are no packet filters, port lock down, or network issues preventing the connection.

What are two reasons the synchronization group is having issues? (Choose two.)

A. Certificates expired on all of the peer LTM devices.

B. Certificates stored for the device trusts on all of the peer LTM devices are corrupted.

C. Admin passwords changed on one of the peer LTM devices that are able to synchronize.

D. Admin password changed on the LTM device NOT receiving the synchronized configurations.

E. Certificates stored for the device trusts on the LTM device NOT receiving the configuration are corrupted.

ANSWER: D E

Explanation:

"Admin password changed on the LTM device NOT receiving the synchronized configurations" is a valid cause because BIG-IP configuration synchronization relies on authenticated communication between members of the device trust. If the administrative credentials expected for communication with the receiving peer no longer match, the sending device cannot successfully authenticate the configuration-sync operation to that member. Since the other members can synchronize, a credential change isolated to the non-receiving device directly fits the observed symptom.

"Certificates stored for the device trusts on the LTM device NOT receiving the configuration are corrupted" is also a valid cause. A device trust establishes secure, certificate-based relationships among BIG-IP devices, and ConfigSync depends on that trust relationship. Corruption of the local trust certificates on one member prevents that member from validating or participating correctly in trusted communication, even when network reachability and port access are confirmed. The failure would therefore be limited to the affected receiving device while the unaffected peers continue to synchronize normally. F5 documents device trust as the prerequisite foundation for synchronization and describes ConfigSync as a device-group operation built on that trusted relationship. See the F5 device trust documentation and F5 ConfigSync documentation.

QUESTION NO: 9

An LTM Specialist is troubleshooting a problem on an eCommerce website. The user browses the online store using port 80, adding items to the shopping cart. The user then clicks the "Checkout" button on the site, which redirects the user to port 443 for the checkout process. Suddenly, the user's shopping cart is shown as empty. The shopping cart data is stored in memory on the server, and the default source address persistence profile is used on both virtual servers.

What is the issue?

A. The port 80 pool member is deleting the user's session cookie.

B. The port 443 pool member is deleting the user's session cookie.

C. The port 80 and port 443 connections are balanced to the same node.

D. The port 80 and port 443 connections are balanced to different nodes.

ANSWER: D

Explanation:

The port 80 and port 443 connections are balanced to different nodes. Source-address persistence creates a persistence record for a client source IP address within the context of the virtual server that processed the connection. Because HTTP on port 80 and HTTPS on port 443 are handled by separate virtual servers, the default source-address persistence behavior does not automatically make their persistence records shared. Consequently, the initial HTTP request can be persisted to one pool member, while the subsequent HTTPS checkout request can be selected by the HTTPS virtual server and sent to another pool member. Since the application stores the shopping-cart state only in the selected server's local memory, that second server has no knowledge of the cart created during the HTTP session and presents an empty cart. A suitable design must ensure that both virtual servers select the same application instance for the user, such as by configuring persistence appropriately across the related virtual servers, or by storing session/cart state in shared application storage rather than only local server memory. F5 documents the source-address persistence profile and its configuration properties in the BIG-IP source-address persistence API reference. Additional persistence-profile configuration guidance is available in the BIG-IP LTM persistence profiles documentation.

QUESTION NO: 10

An LTM Specialist is setting up a monitor for an HTTP 1.1 server. The response to a GET / is:

HTTP/1.1 302 Moved Temporarily

Location: http://www.example.com/new/location.html

Which send string settings should the LTM Specialist use to force a proper response?

A. GET / HTTP/1.0\r\nHost: host.domain.com\r\nConnection: Close\r\n\r\n

B. GET /new/location.html HTTP/1.1\r\nHost: www.example.com\r\nConnection: Close\r\n\r\n

C. GET / HTTP/1.1\r\nHost: www.example.com/new/location.html\r\nConnection: Close\r\n\r\n

D. GET /new/location.html HTTP/1.1\r\nHost: host.domain.com/new/locations.html\r\nConnection: Close\r\n\r\n

ANSWER: B

Explanation:

`GET /new/location.html HTTP/1.1\r\nHost: www.example.com\r\nConnection: Close\r\n\r\n` is the correct monitor send string because the server's HTTP 302 response supplies the target resource in its `Location` header: `http://www.example.com/new/location.html`. To obtain the intended successful response, the monitor must request that redirected path directly, using `/new/location.html` as the request target.

Because the request uses HTTP/1.1, it must also include a Host header identifying the authority named by the redirected URL. Setting `Host: www.example.com` ensures the request is handled by the appropriate HTTP virtual host, which is especially important when multiple sites share the same listener address. The CRLF sequences terminate each HTTP header line, and the final blank line terminates the header section. `Connection: Close` instructs the server to close the monitor connection after responding, making the monitor transaction self-contained.

F5 documents that HTTP monitor send strings are complete HTTP request data and supports escaped CRLF line endings in monitor requests. See the F5 tmsh HTTP monitor reference and F5's guidance for constructing HTTP monitor requests.

QUESTION NO: 11

-- Exhibit –

```
ltm profile httpclass acct_class {
    app-service none
    defaults-from httpclass
    paths { glob:/accounting }
    pool srv1_http_pool
    redirect none
}
ltm profile httpclass marketing_class {
    app-service none
    defaults-from httpclass
    paths { glob:/marketing }
    pool srv1_http_pool
    redirect none
}
ltm profile httpclass default_class {
    app-service none
    defaults-from httpclass
    pool srv2_http_pool
    redirect none
}
ltm virtual http_vs {
    destination 192.168.1.155:http
    http-class {
        acct_class
        marketing_class
        default_class
    }
    ip-protocol tcp
    mask 255.255.255.255
    pool srv2_http_pool
    profiles {
        http { }
        tcp { }
    }
    snat automap
    vlans-disabled
}
```

-- Exhibit --

Refer to the exhibit.

An LTM Specialist is reviewing the virtual server configuration on an LTM device.

Which two actions should the LTM Specialist perform to minimize the virtual server configuration? (Choose two.)

A. Remove 'snat automap' from the virtual server.

B. Remove the 'http' profile from the virtual server.

C. Remove the 'default_class' from the virtual server.

D. Combine 'acct_class' and 'marketing_class' into one class and update associations on the virtual server.

E. Combine 'marketing_class' and 'default_class' into one class and update associations on the virtual server.

ANSWER: C D

Explanation:

Removing the *default_class* from the virtual server reduces configuration because the virtual server's default processing already provides the required fallback behavior. A separate default classification association is therefore unnecessary when it does not introduce a distinct action beyond the virtual server's normal default handling. Removing redundant associations makes the active configuration easier to inspect and lowers the chance that later changes create overlapping or difficult-to-troubleshoot traffic handling.

Combining *acct_class* and *marketing_class* into one class also minimizes the configuration because the two classifications can be represented as a single consolidated match definition while preserving their common association and resulting traffic treatment. The virtual server then needs only one class association rather than maintaining separate objects for equivalent processing. This follows the general BIG-IP LTM design practice of consolidating policy logic where matching traffic receives the same action, while retaining a clear default path for traffic that does not match the consolidated class. F5 documents virtual-server traffic processing and associated policies in its BIG-IP Local Traffic Manager Implementations guide and provides policy configuration guidance in the F5 LTM configuration documentation.

QUESTION NO: 12

An LTM Specialist configured a virtual server to load balance a custom application. The application works when it is tested from within the firewall but it fails when tested externally. The pool member address is 192.168.200.10:80. A capture from an external client shows:

GET /index.jsp HTTP/1.1

Host: 207.206.201.100

User-Agent: Mozilla/5.0 (Windows NT 6.1; WOW64; rv:15.0) Gecko/20100101 Firefox/15.0.1

Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8

Connection: keep-alive

HTTP/1.1 302 Found

DatE. Wed, 17 Oct 2012 23:09:55 GMT

Server: Apache/2.2.15 (CentOS)

Location: http://192.168.200.10/user/home.jsp

Content-LengtH. 304

Connection: close

What is the solution to this issue?

A. Assign a SNAT pool to the virtual server.

B. Add a Web Acceleration Profile to the virtual server.

C. Configure redirect rewrite option in the HTTP profile.

D. Configure a content filter on the backend web server.

ANSWER: C

Explanation:

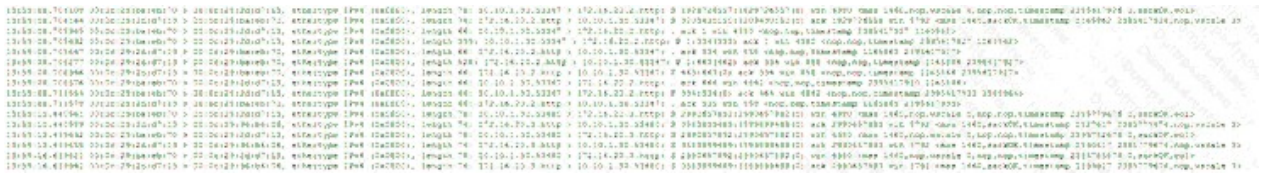
Configure redirect rewrite option in the HTTP profile. The server's HTTP 302 response contains a `Location` header that exposes the pool member's private address, `192.168.200.10`. An external browser follows that redirect and attempts to connect directly to the private address, which is not routable or reachable from outside the firewall. The original request correctly used the public virtual-server address, `207.206.201.100`, so the redirect target must be translated to a client-reachable address as well.

BIG-IP's HTTP profile includes Redirect Rewrite processing specifically for this situation. When enabled with the appropriate rewrite mode, BIG-IP examines applicable HTTP redirect response headers and replaces the internal server address with

the virtual server address presented to the client. The client then follows the rewritten URL back through the virtual server, allowing LTM to select and connect to the pool member while keeping the internal topology undisclosed. This is especially important for applications that generate absolute URLs in 301, 302, 303, or 307 redirects. F5 documents the HTTP profile's redirect-rewrite behavior in its LTM HTTP profile documentation and lists the `redirect-rewrite` setting in the tmsh HTTP profile reference.

QUESTION NO: 13

-- Exhibit –



-- Exhibit --

Refer to the exhibit.

An LTM Specialist configures a virtual server that balances HTTP connections to a pool of three application servers. Approximately one out of every three connections to the virtual server fails.

Which two actions will resolve the problem? (Choose two.)

A. Assign a custom HTTP monitor to the pool.

B. Enable SNAT automap on the virtual server.

C. Verify that port lockdown is set to allow port 80.

D. Verify the default gateway on the application servers.

E. Increase the TCP timeout value in the default TCP profile.

ANSWER: B D

Explanation:

Enabling SNAT automap on the virtual server resolves return-path problems by translating each client source address to an appropriate BIG-IP self IP address. Consequently, pool members send their responses back to BIG-IP, which retains the connection state and can correctly forward each response to its original client. This is especially important when the application servers do not otherwise have a route that returns client-bound traffic through BIG-IP. F5 documents that SNAT automap selects a self IP address for source translation and is commonly used to ensure that return traffic flows through the device.

Verifying the default gateway on the application servers addresses the same requirement at the network-routing layer. Each pool member must have a valid return route for client networks, normally through BIG-IP when it is performing full proxy load balancing. If one application server has an incorrect or missing default gateway, requests directed to that member can fail while requests directed to the other two members succeed, matching the approximately one-in-three failure pattern. Correcting the gateway or an equivalent specific route ensures symmetric traffic flow through BIG-IP and allows the TCP and HTTP connection state to be maintained. See F5's SNAT configuration documentation and F5 routing and traffic-flow training material.

QUESTION NO: 14

Which two items can be logged by the Application Visibility Reporting analytics profile? (Choose two.)

A. User Agent

B. HTTP version

C. HTTP Response Codes

D. Per Virtual Server CPU Utilization

ANSWER: A C

Explanation:

User Agent and **HTTP Response Codes** are analytics data items that an Application Visibility Reporting analytics profile can collect and log for HTTP traffic. User-Agent information identifies the client software associated with requests, allowing Analytics reports to characterize the browsers, operating systems, and other client types accessing an application. This is useful for recognizing client populations and investigating whether a traffic or performance pattern is associated with a particular client type.

HTTP response-code logging records the status returned by the application or virtual server for transactions. Analytics can use these response values in reporting to show application outcomes and trends, such as successful responses and error responses. Capturing this transaction-level HTTP information is central to the Analytics profile's role of providing visibility into application request and response behavior. The BIG-IP Analytics documentation describes Analytics as collecting HTTP transaction data for reports and identifies client details and HTTP response information among the reported metrics. See the F5 BIG-IP Analytics overview and the F5 documentation for configuring Analytics profiles.

QUESTION NO: 15

An LTM Specialist configures two LTM devices in a high-availability pair with trusts established and device groups configured properly using network failover. After several months, the LTM Specialist notices that changes made to one LTM device do NOT cause the synchronization status to update to "changes pending," and this device does NOT synchronize with the device group.

Which two steps should the LTM Specialist take to identify the issue? (Choose two.)

A. Verify that NTP is synchronized.

B. Verify the network connectivity between the devices.

C. Verify that the devices are not using self-signed certificates.

D. Verify that ConfigSync is using the management IP address.

E. Verify that port lockdown on the ConfigSync interface is set to allow port 1026.

ANSWER: A B

Explanation:

Verify that NTP is synchronized. Accurate, synchronized system time is foundational to the device trust and ConfigSync relationship because the BIG-IP devices use certificates and secure communications to authenticate one another. If the clocks drift substantially over time, certificate-validity checks and trust-related communications can fail or behave inconsistently. Checking NTP status on both peers is therefore an appropriate first diagnostic action when an established high-availability pair later stops recognizing and exchanging configuration changes. F5 recommends configuring reliable NTP sources for BIG-IP systems, particularly where device trust and high availability are in use.

Verify the network connectivity between the devices. ConfigSync requires reliable bidirectional connectivity between the addresses selected for synchronization, while network failover also requires communication between the peers. Test reachability using the intended self or management addresses, confirm routing and VLAN availability, and ensure that intervening firewalls or security devices permit the required BIG-IP device-service-clustering traffic. Loss of this communication can prevent a peer from receiving configuration updates and from maintaining a current synchronization state. F5 documents the device-service-clustering communication requirements and the role of network connectivity in ConfigSync and failover operation in the BIG-IP Device Service Clustering documentation and its BIG-IP system ports and protocols reference.

QUESTION NO: 16

An LTM Specialist receives a request to monitor the network path through a member, but NOT the member itself.

Which monitor option should the LTM Specialist enable or configure?

A. Reverse

B. Up interval

C. Transparent

D. Alias address

E. Time until up

ANSWER: C

Explanation:

Transparent is the appropriate monitor setting when the goal is to test connectivity through a pool member to another destination, rather than to assess the health of the pool member itself. With transparent monitoring enabled, BIG-IP sends the monitor probe to the pool member's address, but uses the monitor's alias address as the intended destination. The pool member therefore forwards the probe toward that destination, allowing the monitor result to reflect whether the required network path and remote endpoint are reachable through the member. This is useful when the member acts as a router, firewall, gateway, or other transit device and the service owner needs validation of the path beyond it. Configure an appropriate *alias address* for the remote endpoint together with transparent monitoring; the alias identifies where the probe is logically directed, while transparent mode causes BIG-IP to send it via the pool member. F5 documents the `transparent` monitor property as enabling monitoring through a pool member and describes its relationship to the alias address. See the F5 tmsh ltm monitor http reference and the LTM monitors reference.

QUESTION NO: 17

An LTM Specialist needs to prevent a pool member from receiving new load-balanced connections while allowing existing connections to complete.

Which two pool member session states can be used? (Choose two.)

A. Disabled

B. Forced Offline

C. Available

D. Checking

E. Unknown

ANSWER: A B

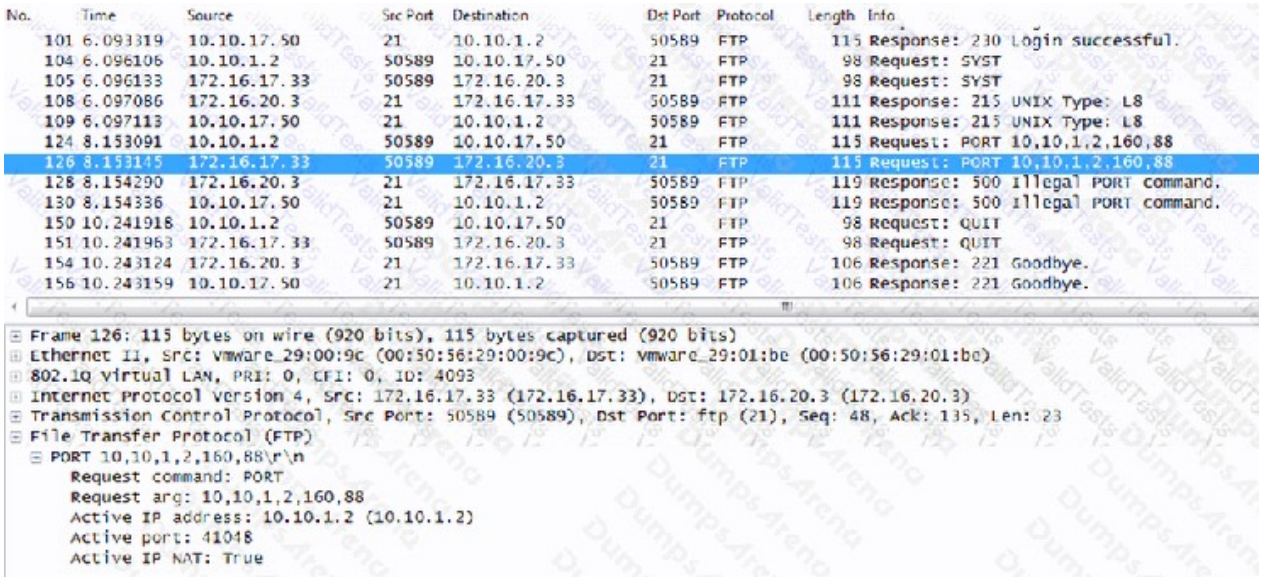
Explanation:

Disabled can be used to remove a pool member from normal new load-balancing decisions while allowing existing connections to continue. This is useful when an administrator needs to drain traffic before performing maintenance.

Forced Offline also prevents BIG-IP from selecting the member for new connections while allowing established connections to complete. It provides a clear administrative state indicating that the member has intentionally been removed from service. The appropriate selection depends on the operational procedure and desired displayed state, but both states stop normal new connection selection. See the F5 tmsh pool reference.

QUESTION NO: 18

-- Exhibit –



-- Exhibit --

Refer to the exhibit.

An LTM Specialist is investigating reports that users are unable to perform some commands through an FTP virtual server. The users are receiving the FTP error "500 Illegal PORT command." The virtual server is configured to SNAT using automap. The LTM Specialist performs a capture on the server side of the LTM device.

Why is the server returning this error?

A. LIST command disallowed

B. PORT command disallowed

C. Active IP address in PORT command

D. Active IP address in LOGIN command

ANSWER: C

Explanation:

Active IP address in PORT command is correct because FTP active mode embeds the client's listening address and port in the clear-text PORT control-channel command. With SNAT automap, the BIG-IP system changes the source address of the client-to-server control connection to a self IP address. If the PORT command still carries the client's original active address, the server observes an address in the command that does not match the source of the established control connection. FTP servers commonly reject that condition as an illegal PORT command because accepting it could permit FTP bounce behavior, in which a client asks a server to initiate a data connection to an unrelated host. The server therefore returns the 500 response before establishing the active-mode data connection. BIG-IP's FTP profile is designed to inspect FTP control traffic and translate embedded addressing information when required, maintaining consistency between the SNATted control connection and the FTP data-channel negotiation. This behavior follows the FTP PORT command definition in RFC 959. For BIG-IP profile behavior and FTP protocol handling, see the BIG-IP Local Traffic Management Profiles Reference.

QUESTION NO: 19

-- Exhibit –

```
GET / HTTP/1.1
Host: www.example.com
User-Agent: Mozilla/5.0 (Windows NT 6.1; rv:16.0) Gecko/20100101 Firefox/16.0
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-US,en;q=0.5
Accept-Encoding: gzip, deflate
DNT: 1
Connection: keep-alive


HTTP/1.1 302 Moved Temporarily
Content-Length: 0
Location: https://www.example.com
Date: Tue, 23 Oct 2012 18:05:57 GMT
Server: Apache/2.2.22 (FreeBSD) PHP/5.4.4 mod_ssl/2.2.22 OpenSSL/0.9.8q DAV/2
Accept-Ranges: bytes
Connection: Keep-Alive
Content-Type: text/html
Set-Cookie: sessionid=a4531785-7012-46aa-b5fe-a54be482b61a; path=/
```

-- Exhibit --

Refer to the exhibit.

An LTM Specialist is performing an HTTP trace on the client side of the LTM device and notices there are many undesired headers being sent by the server in the response. The LTM Specialist wants to remove all response headers except "Set-Cookie" and "Location."

How should the LTM Specialist modify the HTTP profile to remove undesired headers from the HTTP response?

A. Enter the desired header names in the 'Request Header Insert' field.

B. Enter the undesired header names in the 'Request Header Erase' field.

C. Enter the undesired header names in the 'Response Header Erase' field.

D. Enter the desired header names in the 'Response Headers Allowed' field.

ANSWER: D

Explanation:

Enter the desired header names in the 'Response Headers Allowed' field. This HTTP profile setting is an allowlist for response headers: BIG-IP retains and forwards only the response header names specified in the field, removing other response headers before the response is delivered to the client. Therefore, configuring `Set-Cookie` and `Location` as the allowed response headers achieves the stated requirement of removing every other server-supplied response header while preserving those two headers. This is particularly useful when a policy requires strict control over the metadata exposed to clients, rather than selectively removing a known and finite list of unwanted headers. The setting applies to HTTP responses, which is essential because the observed undesired headers are travelling from the server through the LTM device to the client. F5 documents the HTTP profile's response-header controls, including the `response-headers-allowed` property, in its TMSH reference. The applicable profile must be associated with the virtual server handling the traffic for the filtering behavior to take effect. See the F5 TMSH HTTP profile reference and the F5 HTTP profile documentation.

QUESTION NO: 20

An HTTPS virtual server must authenticate clients by validating the certificate presented during the TLS handshake.

Which two Client SSL profile settings should the LTM Specialist configure? (Choose two.)

A. Client Certificate set to require

B. Trusted Certificate Authorities

C. Server Certificate

D. SSL Forward Proxy

E. Server Name

ANSWER: A B

Explanation:

Client Certificate set to **require** instructs BIG-IP to request a client certificate and to reject the TLS handshake when the client does not provide an acceptable certificate. This enforces mutual TLS authentication rather than merely allowing optional client certificate presentation.

Trusted Certificate Authorities identifies the certificate bundle BIG-IP uses to validate the issuer chain of the client certificate. BIG-IP must trust the CA that issued the client certificate before it can successfully authenticate the client. The Client SSL profile provides both settings for client-certificate authentication. See the F5 tmsh Client SSL profile reference.

QUESTION NO: 21

What does the following iRule do?

when CLIENT_ACCEPTED {

if { [matchclass [IP::client_addr] equals WebClient1-Whitelist1] }{

#log local0. "Valid client IP: [IP::client_addr] - forwarding traffic"

#Pool WebClient1

} else {

log local0. "Invalid client IP: [IP::client_addr] - discarding"

discard

}

}

A. The iRule compares a client IP to a list. If the client IP is on the list, discard and log the discard.

B. The iRule compares a client IP to a list. If the client IP is NOT on the list, discard and log the discard.

C. The iRule compares a client IP to a list. If the client IP is on the list, the client is sent to Pool WebClient1. Otherwise, discard and log the discard.

D. The iRule compares a client IP to a list. If the client IP is NOT on the list, the client is sent to Pool WebClient1. Otherwise, discard and log the discard.

ANSWER: B

Explanation:

The iRule compares the connecting client's source address, obtained through `IP::client_addr`, against the data group or list named `WebClient1-Whitelist1`. It runs during the `CLIENT_ACCEPTED` event, which occurs when BIG-IP accepts a client-side connection. If the source address matches an entry in that list, the conditional branch contains only commented-out commands, so the iRule takes no traffic-altering action and normal virtual-server processing continues. In particular, `#Pool WebClient1` is a comment, not an active `pool` command, and therefore does not direct the connection to that pool.

When the client address does not match the whitelist, the `else` branch writes an "Invalid client IP" message to the local logging facility and invokes `discard`. The `discard` command silently drops the connection/traffic, preventing the non-whitelisted client from proceeding through the virtual server. Thus, the effective behavior is to allow normal processing for

listed client IP addresses and log then discard clients that are not listed. F5 documents the `matchclass` command's use for matching values against data groups and documents `discard` as silently dropping traffic: matchclass iRules reference and discard iRules reference.

QUESTION NO: 22

A new VLAN vlan301 has been configured on a highly available LTM device in partition ApplicationA. A new directly connected backend server has been placed on vlan301. However, there are connectivity issues pinging the default gateway. The VLAN self IPs configured on the LTM devices are 192.168.0.251 and 192.168.0.252 with floating IP 192.168.0.253. The LTM Specialist needs to perform a packet capture to assist with troubleshooting the connectivity.

Which command should the LTM Specialist execute on the LTM device command line interface to capture the attempted pings to the LTM device default gateway on VLAN vlan301?

A. tcpdump -ni /ApplicationA/vlan301 'host 192.168.0.253'

B. tcpdump -ni vlan301 'host 192.168.0.253'

C. tcpdump -ni /ApplicationA/vlan301 'host 192.168.0.251 or host 192.168.0.252'

D. tcpdump -ni vlan301 'host 192.168.0.251 or host 192.168.0.252'

ANSWER: A

Explanation:

`tcpdump -ni /ApplicationA/vlan301 'host 192.168.0.253'` is the appropriate capture command because the default gateway presented to hosts on this highly available VLAN is the traffic-group floating self IP, 192.168.0.253. A host pinging its configured gateway sends ICMP traffic to that floating address, so filtering on the floating IP isolates the relevant request and reply packets while avoiding unrelated VLAN traffic. The `-i` argument identifies the capture interface, and BIG-IP VLAN objects that reside outside the Common partition should be referenced with their fully qualified partition path. Therefore, `/ApplicationA/vlan301` correctly identifies the VLAN configured in the ApplicationA partition. The `-n` argument prevents name resolution during the capture, which keeps output immediate and avoids DNS-related noise or delays. This capture can show whether the echo requests arrive at BIG-IP on vlan301 and whether echo replies leave through that VLAN, providing the essential packet-level evidence for investigating the gateway-connectivity symptom. F5 documents VLAN object naming and partitioned configuration objects in the tmsh net vlan reference; its packet-capture guidance is available in F5 article K411.

QUESTION NO: 23

Which command will identify the active LTM device currently handling client traffic?

A. b ha table show

B. tmsh list /sys ha-status

C. tmsh show /cm traffic-group

D. tmsh run /sys failover standby

E. tmsh show /sys ha-status all-properties

ANSWER: C

Explanation:

`tmsh show /cm traffic-group` is the correct command because traffic groups are the BIG-IP high-availability objects that own floating traffic-management resources, including floating self IP addresses and virtual addresses. Client traffic reaches the device that is currently active for the traffic group containing those resources. The command displays traffic-group runtime status and identifies the device on which each traffic group is active, allowing an administrator to determine

which LTM peer is presently servicing traffic for the associated virtual servers. This is particularly useful in device-service clustering configurations because different traffic groups can be active on different devices when an active-active design is used; therefore, a general local high-availability state alone may not establish ownership for a specific set of traffic resources. F5 documents traffic groups as the mechanism used to control which device is active for floating objects and provides the `show cm traffic-group` command for viewing their status. See the F5 tmsh cm traffic-group reference and the F5 traffic groups documentation.

QUESTION NO: 24

An FTP monitor is NOT working correctly.

Which three pieces of information does the LTM Specialist need to provide to ensure a properly working FTP monitor? (Choose three.)

A. alias

B. File path

C. username

D. password

E. FTP server port

F. FTP server IP address

ANSWER: B C D

Explanation:

An FTP monitor must be able to authenticate to the FTP service and perform the configured file check. The **username** provides the account name sent in the FTP USER command, and the **password** supplies the corresponding credential for the FTP PASS command. Supplying valid credentials is essential when the server does not permit anonymous access or when the monitored account has restricted access.

The **File path** identifies the file resource the monitor uses to validate the FTP service. In BIG-IP FTP monitor configuration, this is represented through the directory and filename settings. The monitor can change to the configured directory and retrieve or check the specified file, confirming that the FTP server is not merely reachable but can complete the required authenticated FTP operation. The monitored pool member defines the service destination used by the monitor, while the FTP monitor settings define the login and file-validation transaction.

For the FTP monitor's configurable properties, including username, password, directory, and filename, see the F5 tmsh reference for ltm monitor ftp. F5's monitor overview and configuration guidance is also available in the BIG-IP health monitor documentation.

QUESTION NO: 25

-- Exhibit –

-- Exhibit --

Refer to the exhibit.

A pair of LTM devices are configured for HA. The LTM Specialist observes from a capture that there is a successful connection from a client directly to a web server and an unsuccessful connection from a client via the LTM device to the same web server.

Which two solutions will solve the configuration problem? (Choose two.)

A. Configure SNAT on the pool.

B. Configure SNAT on the virtual server.

C. Change server default gateway to point at LTM internal self IP.

D. Change server default gateway to point at LTM internal floating IP.

ANSWER: B D

Explanation:

Configure SNAT on the virtual server and change the server default gateway to point at LTM internal floating IP. These changes address the return-path requirement for a full-proxy virtual server. With source network address translation enabled on the virtual server, the pool member receives traffic with an LTM-controlled source address. Its response is therefore directed back to the LTM, allowing the system to match the response to the server-side flow and forward it through the established client-side connection. SNAT is especially appropriate where the server does not otherwise have a route that reliably returns client traffic through the traffic-management device.

Using the internal *floating* self IP as the server's default gateway provides the highly available routing endpoint. The traffic group owns that floating address, so it moves to the peer device when a failover occurs. Servers can continue to use the same default gateway address without requiring a routing change, while the active device remains able to receive server return traffic. F5 describes SNAT behavior and its role in ensuring return traffic passes through the BIG-IP system in its SNAT overview, and documents floating self IP addresses and traffic-group failover in the Device Service Clustering documentation.

QUESTION NO: 26

A source address persistence profile is configured with a mask of 255.255.255.0.

Which clients will use the same persistence record?

A. All clients with the same first 24 bits of the source address

B. All clients with the same first 16 bits of the source address

C. Only clients with an identical source address

D. All clients using the same source port

ANSWER: A

Explanation:

Clients whose source addresses are in the same **/24 network** use the same persistence record. The mask 255.255.255.0 preserves the first 24 bits of the client address when BIG-IP creates and looks up the source-address persistence key. For example, clients 192.0.2.10 and 192.0.2.200 are treated as the same persistence source, while 192.0.3.10 uses a different record.

A source address persistence mask determines the portion of the client address used for persistence matching. A /32 mask creates a record per host, while a broader mask groups more client addresses into the same persistence scope. See the F5 tmsh source address persistence profile reference.

QUESTION NO: 27

The output of a tmsh command is: ---------- Net::Interface Name Status Bits Bits Errs Errs Drops Drops Colli In Out In Out In Out sions ---------- 1.1 down 0 0 0 0 0 0 0 1.2 up 191.4K 0 0 0 374 0 0 1.3 down 0 0 0 0 0 0 0 1.4 up 22.5K 0 0 0 44 0 0 2.1 miss 0 0 0 0 0 0 0 2.2 miss 0 0 0 0 0 0 0 mgmt up 43.2G 160.0G 0 0 0 0 0

Which command was executed on the LTM device to show the output?

A. tmsh show /net interface

B. tmsh /net show interface status

C. tmsh /net show interface

D. tmsh show /net interface status

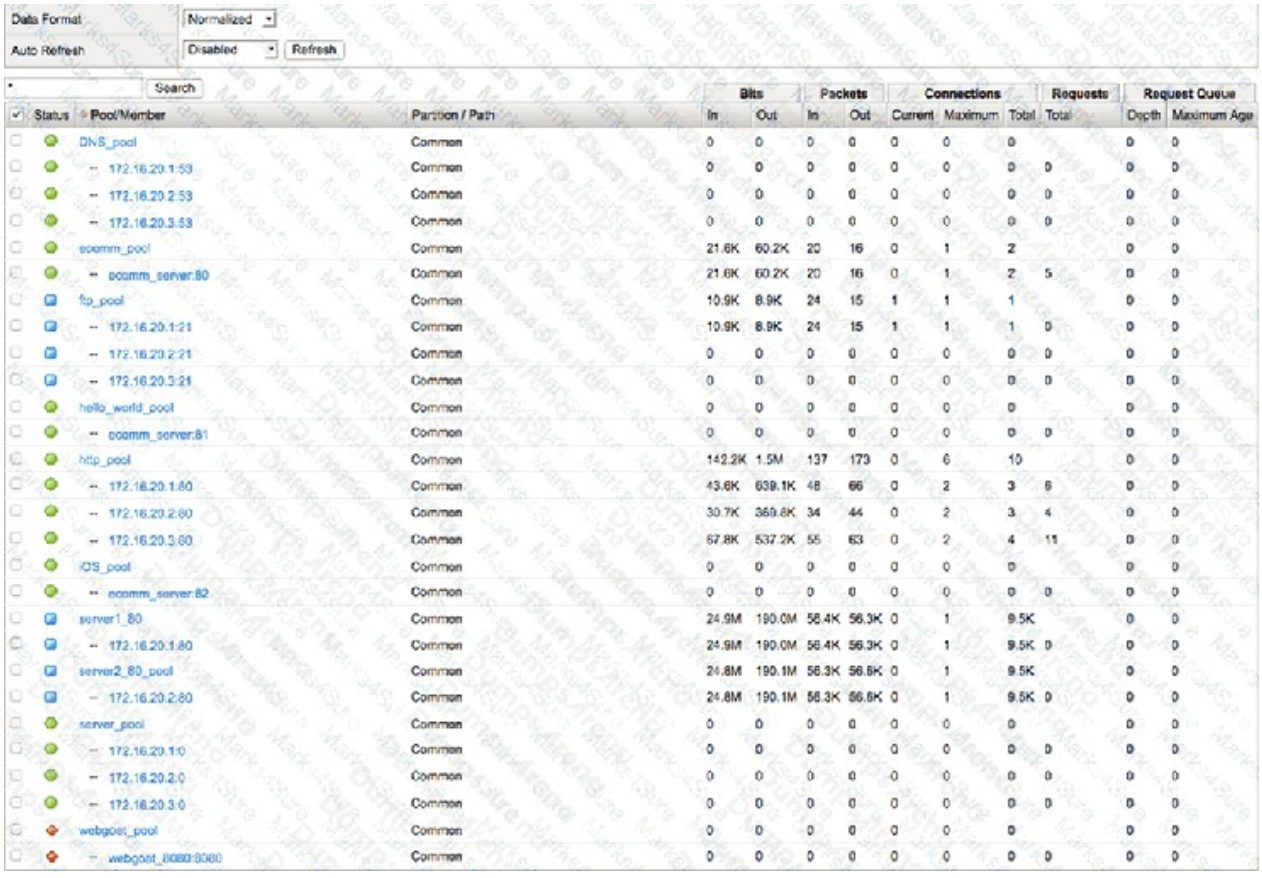
ANSWER: A

Explanation:

`tmsh show /net interface` is the command used to display operational status and cumulative traffic/error counters for BIG-IP network interfaces. Its output is formatted as the `Net::Interface` table and includes each interface name, link status, inbound and outbound bit counts, inbound and outbound error counts, drop counts, and collision counts. This directly matches the displayed fields, including data interfaces such as `1.1` and `1.2` as well as the dedicated `mgmt` interface. The status values shown in this operational view can include `up`, `down`, and `miss`, allowing an administrator to quickly identify whether an interface is available and whether it is experiencing packet drops or physical-layer errors. The leading slash identifies the `net` module path; it is an accepted fully qualified tmsh object path. F5 documents the `show net interface` command as the interface statistics and status display command, with interface counters presented as operational data rather than configuration details. See the F5 tmsh reference for net interface and the F5 guide to tmsh command syntax.

QUESTION NO: 28

-- Exhibit –



-- Exhibit --

Refer to the exhibit.

Which pool can be removed without affecting client traffic?

A. ftp_pool

B. http_pool

C. server1_80

D. server_pool

ANSWER: D

Explanation:

server_pool can be removed without affecting client traffic because the exhibit's object relationships show that it is not selected by any client-facing virtual server, nor is it reached through another active traffic-processing object. In BIG-IP LTM, a pool is simply a collection of members that becomes relevant to client connections only when a virtual server, policy, iRule, or another configured traffic workflow directs connections to it. A pool that has no such active reference is an unused configuration object; deleting it does not alter virtual-server matching, connection handling, load-balancing decisions, or the destination selected for current client traffic.

This distinction is important during configuration cleanup: pool names alone do not establish traffic dependency. The relevant check is whether a virtual server's *pool* property—or another applicable traffic rule—references the pool. BIG-IP documentation describes a virtual server as the listener that receives client traffic and identifies the pool to which matching connections are sent. The pool object then distributes those connections to its configured members. Since **server_pool** has no active traffic reference in the displayed configuration, it is safe to remove from the standpoint of client traffic. See F5's Virtual Servers and Pools overview and the BIG-IP pool API reference.

QUESTION NO: 29

-- Exhibit –



-- Exhibit --

Refer to the exhibit.

A pair of LTM devices are configured for HA. The LTM Specialist observes from a capture that there is a successful connection from a client directly to a web server and an unsuccessful connection from a client via the LTM device to the same web server.

Which two solutions will solve the configuration problem? (Choose two.)

A. Configure SNAT on the pool.

B. Configure SNAT on the virtual server.

C. Change server default gateway to point at LTM internal self IP.

D. Change server default gateway to point at LTM internal floating IP.

ANSWER: B D

Explanation:

Configuring SNAT on the virtual server solves the asymmetric return-path problem by translating the client source address to an address owned by the BIG-IP system. The web server then returns its response to BIG-IP, which has connection state for the flow and can translate and forward the response back to the original client. This is particularly appropriate when the server's existing default gateway does not route client networks through BIG-IP.

Changing the server default gateway to point at the LTM internal floating IP also ensures that responses for connections handled by either device in the HA pair return through the active traffic-processing unit. A floating self IP is the correct highly available gateway address because it follows the active device after failover. Consequently, server return traffic continues to traverse BIG-IP, allowing the full-proxy connection to complete and preserving service availability during an HA state transition.

F5 documents SNAT as a mechanism that translates source addresses for virtual-server traffic and describes floating self IP addresses as the appropriate shared traffic addresses in device service clustering. See the F5 TMSH SNAT reference and F5 Device Service Clustering documentation.

QUESTION NO: 30

The LTM device is configured to provide load balancing to a set of web servers that implement access control lists (ACL) based on the source IP address of the client. The ACL is at the network level and the web server is configured to send a TCP reset back to the client if it is NOT permitted to connect.

The virtual server is configured with the default OneConnect profile.

The ACL is defined on the web server as:

Permit: 192.168.136.0/24

Deny: 192.168.116.0/24

The packet capture is taken of two individual client flows to a virtual server with IP address 192.168.136.100.

Client A - Src IP 192.168.136.1 - Virtual Server 192.168.136.100:

Clientside:

09:35:11.073623 IP 192.168.136.1.55684 > 192.168.136.100.80: S 869998901:869998901(0) win 8192

09:35:11.073931 IP 192.168.136.100.80 > 192.168.136.1.55684: S 2273668949:2273668949(0) ack 869998902 win 4380

09:35:11.074928 IP 192.168.136.1.55684 > 192.168.136.100.80: . ack 1 win 16425

09:35:11.080936 IP 192.168.136.1.55684 > 192.168.136.100.80: P 1:299(298) ack 1 win 16425

09:35:11.081029 IP 192.168.136.100.80 > 192.168.136.1.55684: . ack 299 win 4678

Serverside:

09:35:11.081022 IP 192.168.136.1.55684 > 192.168.116.128.80: S 685865802:685865802(0) win 4380

09:35:11.081928 IP 192.168.116.128.80 > 192.168.136.1.55684: S 4193259095:4193259095(0) ack 685865803 win 5840

09:35:11.081943 IP 192.168.136.1.55684 > 192.168.116.128.80: . ack 1 win 4380

09:35:11.081955 IP 192.168.136.1.55684 > 192.168.116.128.80: P 1:299(298) ack 1 win 4380

09:35:11.083765 IP 192.168.116.128.80 > 192.168.136.1.55684: . ack 299 win 108

Client B - Src IP 192.168.116.1 - Virtual Server 192.168.136.100:

Clientside:

09:36:11.244040 IP 192.168.116.1.55769 > 192.168.136.100.80: S 3320618938:3320618938(0) win 8192

09:36:11.244152 IP 192.168.136.100.80 > 192.168.116.1.55769: S 3878120666:3878120666(0) ack 3320618939 win 4380

09:36:11.244839 IP 192.168.116.1.55769 > 192.168.136.100.80: . ack 1 win 16425

09:36:11.245830 IP 192.168.116.1.55769 > 192.168.136.100.80: P 1:299(298) ack 1 win 16425

09:36:11.245922 IP 192.168.136.100.80 > 192.168.116.1.55769: . ack 299 win 4678

Serverside:

09:36:11.245940 IP 192.168.136.1.55684 > 192.168.116.128.80: P 599:897(298) ack 4525 win 8904

09:36:11.247847 IP 192.168.116.128.80 > 192.168.136.1.55684: P 4525:5001(476) ack 897 win 142

Why was the second client flow permitted by the web server?

A. A global SNAT is defined.

B. SNAT automap was enabled on the virtual server.

C. The idle TCP session from the first client was re-used.

D. A source address persistence profile is assigned to the virtual server.

ANSWER: C

Explanation:

The idle TCP session from the first client was re-used. OneConnect enables BIG-IP to detach eligible HTTP requests from clientside connections and multiplex them over established, idle serverside TCP connections. With the default OneConnect profile, the source mask permits a broadly shared serverside connection pool. In this capture, the second client establishes its own clientside TCP connection to the virtual server, but its request is sent serverside on the existing connection from 192.168.136.1:55684 to 192.168.116.128:80. The serverside packets for the later request retain that same source IP address and source port, and their TCP sequence numbers continue the already-established session rather than showing a new TCP three-way handshake.

Consequently, the web server sees the request on a TCP connection that was originally established from 192.168.136.1, an address in the permitted subnet. Because the network ACL is evaluated when the TCP connection is established, it does not see a new serverside connection sourced from 192.168.116.1 for the second request. BIG-IP therefore reuses the authorized connection even though the new clientside flow originates from the denied subnet. F5 documents that OneConnect supports reuse of serverside connections for multiple client requests; its source mask controls which client connections can share them. See the F5 OneConnect profile reference and F5 OneConnect implementation documentation.

QUESTION NO: 31

An LTM Specialist is running the following packet capture on an LTM device:

ssldump -Aed -ni vlan301 'port 443'

Which two SSL record message details will the ssldump utility display by default? (Choose two.)

A. HTTP Version

B. User-Agent

C. ClientHello

D. ServerHello

E. Issuer

ANSWER: C D

Explanation:

`ClientHello` and `ServerHello` are TLS handshake messages that `ssldump` decodes and reports as part of its SSL/TLS protocol output. The ClientHello initiates TLS negotiation and contains such protocol-level information as the offered TLS

version, cipher suites, extensions, and potentially the server name indication. The ServerHello is the corresponding server response that selects the negotiated parameters. Because these are visible handshake records, they can be displayed without requiring decryption of the protected application payload.

In the supplied command, `-i vlan301` selects the BIG-IP interface/VLAN capture point, `-n` suppresses hostname resolution, and the Berkeley Packet Filter expression restricts capture to TCP port 443. The SSL/TLS decoding and record-display flags make the handshake exchange useful for troubleshooting negotiation failures, such as protocol-version or cipher-suite mismatch. The `ssldump` documentation identifies ClientHello and ServerHello as decoded SSL handshake protocol messages; see the ssldump manual page and the TLS handshake protocol specification.

QUESTION NO: 32

The active LTM device in a high-availability (HA) pair performs a failover at the same time the network team reports an outage of a switch on the network.

Which two items could have caused the failover event? (Choose two.)

A. a VLAN fail-safe setting

B. a monitor on a pool in an HA group

C. the standby LTM that was rebooted

D. an Auditor role that has access to the GUI

E. the standby LTM that lost connectivity on the failover VLAN

ANSWER: A B

Explanation:

A VLAN fail-safe setting can cause a device to fail over when the BIG-IP system detects no traffic on the configured VLAN for the configured fail-safe timeout. A switch outage can interrupt traffic to that VLAN, allowing the fail-safe timer to expire. When the VLAN is configured with the appropriate fail-safe action, this condition can initiate a traffic-group failover to the peer device. F5 documents the VLAN fail-safe timeout and action as properties used to detect a VLAN failure and determine the system response.

A monitor on a pool in an HA group can also cause failover. An HA group assigns scores to resources, including pools, and the score is affected by the monitored availability of pool members. If a switch outage makes pool members unreachable, the pool can lose its HA-group score. When the active device's resulting HA-group score falls below the configured active threshold, the device can fail over the traffic group to the peer. This provides service-aware failover based on application-resource health, rather than relying solely on device or heartbeat connectivity. See F5's net vlan reference for fail-safe properties and its sys ha-group reference for HA-group scoring and thresholds.