

DUMPS ARENA

BIG-IP ASM Specialist

F5 303

Version Demo

Total Demo Questions: 54

Total Premium Questions: 540

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Application Security	210
Topic 2, Advanced Application Security	6
Topic 3, Troubleshooting	293
Topic 4, Mix Questions	31
Total	540

QUESTION NO: 1

An LTM Specialist needs to create a virtual server to pass TCP traffic to three pool members.

Which two virtual server types should be used to meet the requirements? (Choose two)

- A. Performance (Layer 4)
- B. Standard
- C. Forwarding (IP)
- D. Stateless
- E. Forwarding (Layer 2)

ANSWER: A B

Explanation:

Performance (Layer 4) and **Standard** virtual servers both support TCP connections and can distribute those connections among the members of an associated pool. A Standard virtual server is the general-purpose full-proxy option in BIG-IP LTM. It creates a client-side and server-side flow, applies the configured TCP profile and other traffic-processing features, selects an eligible pool member using the pool's load-balancing method, and forwards the TCP connection to that member. This is the normal choice when full LTM functionality is required.

A Performance (Layer 4) virtual server is also suitable when the requirement is specifically to pass and load balance TCP traffic, but with a lower-overhead Layer 4 processing model. It can use a pool and its load-balancing method to select from the three members while handling TCP traffic efficiently. The choice between these two types depends on whether the deployment needs the broader full-proxy features of a Standard virtual server or the streamlined Layer 4 handling of a Performance (Layer 4) virtual server. F5's virtual-server documentation describes these virtual server types and their traffic-processing behavior: [BIG-IP LTM virtual servers](#) and [F5 virtual server types overview](#).

QUESTION NO: 2

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, returned by `IP::client_addr`, against the data group named `WebClient1-Whitelist1`. It runs at the `CLIENT_ACCEPTED` event, which occurs when BIG-IP accepts a client-side connection. The `matchclass ... equals` expression determines whether the client IP address has an exact matching entry in that whitelist data group. When there is no matching entry, execution enters the `else` block. BIG-IP writes an "Invalid client IP" message, including the client address, to the local `local0` log facility and then executes `discard`. The `discard` command silently drops the connection/traffic rather than forwarding it. Therefore, the effective behavior is to permit connections from listed addresses to continue through normal virtual-server processing while discarding and logging connections from addresses not on the list. Although the `valid-client` log statement and `pool-selection` statement appear in the source, both lines begin with `#`, so they are comments and are not executed. F5 documents the `CLIENT_ACCEPTED` event at [F5 iRules CLIENT_ACCEPTED reference](#) and the silent-drop behavior at [F5 iRules discard reference](#).

QUESTION NO: 3

An ASM event log shows that a request was blocked by an attack signature. The application team believes the event might be a false positive, but the security team must determine whether the same signature should remain enforced elsewhere in the application.

Which two event details should the ASM Specialist review first? (Choose two.)

- A. The matched attack signature ID
- B. The pool load-balancing method
- C. The VLAN configuration of the virtual server
- D. The system management route
- E. The request URL and the location in which the signature matched

ANSWER: A E

Explanation:

The matched attack signature ID identifies the precise signature that caused the violation and allows the specialist to assess its purpose, accuracy, and other events associated with it. The request URL and matched location identify where the signature matched, which is necessary when deciding whether any exception should be narrowly scoped to a specific URL, parameter, or other request element.

The pool load-balancing method, VLAN configuration, and system management route do not establish why an ASM attack signature matched a request.

QUESTION NO: 4

Users are experiencing low throughput when downloading large files over a high-speed WAN connection. Extensive packet loss was found to be an issue but CANNOT be eliminated.

Which two TCP profile settings should be modified to compensate for the packet loss in the network? (Choose two.)

- A. slow start
- B. proxy options
- C. proxy buffer low
- D. proxy buffer high
- E. Nagle's algorithm

ANSWER: C D

Explanation:

proxy buffer low and **proxy buffer high** are the appropriate TCP profile settings because BIG-IP uses the proxy-buffer watermarks to control flow between the client-side and server-side TCP connections that it independently proxies. On a high-bandwidth WAN with unavoidable loss, one side of the proxied connection can temporarily deliver data more quickly than the other side can successfully transmit and recover from dropped packets. Adequate proxy-buffer capacity lets BIG-IP retain more data during that recovery interval instead of applying flow control prematurely. This supports sustained transfer of large files while TCP retransmission and congestion recovery occur across the lossy path.

The low watermark determines when BIG-IP can resume reading data after buffer use has fallen, while the high watermark determines when it pauses reads because buffered data has reached the configured threshold. Adjusting these values together provides a usable buffering range: the high threshold accommodates the temporary backlog caused by packet loss, and the low threshold governs when traffic intake resumes. The values must be sized with available BIG-IP memory and expected concurrent flows in mind. F5 documents these controls as TCP profile properties in the [tmsh TCP profile reference](#); F5's [TCP profile settings reference](#) also describes proxy buffering behavior.

QUESTION NO: 5

An LTM Specialist observes decreased performance and intermittent connection reap LTM system.

Based on the configuration, which action will address these issues?

- A. Use an optimized TCP profile.
- B. Use a FastL4 profile on the virtual server
- C. Use a default caching profile on the virtual server.
- D. Use a shorter idle timeout on the TCP profile.

ANSWER: D

Explanation:

Use a shorter idle timeout on the TCP profile. BIG-IP maintains connection-flow state for TCP connections until the configured idle timeout expires. When a deployment has many inactive or abandoned connections, retaining that state for too long consumes connection-table capacity and, where source-address translation is in use, can also retain translation mappings and ephemeral ports. This can degrade throughput and lead to intermittent connection handling problems as available system resources become constrained. Reducing the TCP idle timeout causes inactive flows to be removed sooner, releasing their associated resources so that active and new client connections can be processed more consistently. The value should still accommodate the application's expected idle periods; it is a capacity and lifecycle tuning setting, not a replacement for application keepalives. F5 documents the TCP profile's idle-timeout setting as the period that a connection may remain idle before BIG-IP closes it, and documents that connection and translation state are maintained by the traffic-management system while flows exist. See the [F5 TCP::idle timeout reference](#) and the [BIG-IP LTM TCP profile documentation](#).

QUESTION NO: 6

A BIG-IP system has the following configuration:

- * SNAT is set to Auto Map
- * There are two VLANs internal and external
- * Default route is pointed to the gateway on external VLAN
- * Self P for internal VLAN is 1921.1.2
- * Self IP for external VLAN is 192.1.2.2

- * Floating IP addresses for internal VLAN is 192.1.1.1
- * Floating IP addresses for external VLAN is 192.1.2.1
- * The Virtual Server IP address is 192.1.1.100

Which IP address does the BIG-IP system use first when traffic reaches the servers on the internal VLAN?

- A. 192.1.1.100
- B. 192.1.2.2
- C. 192.1.1.1
- D. 192.1.2.1

ANSWER: C

Explanation:

192.1.1.1 is the address BIG-IP uses first. With SNAT set to Auto Map, BIG-IP selects a self IP address associated with the VLAN through which it sends the translated connection toward the destination server. Because the pool members are reached through the internal VLAN, the source address selected for the server-side flow must be an internal-VLAN self IP address.

When both a floating and a non-floating self IP are available on the applicable egress VLAN, Auto Map prefers the floating self IP. This preserves connection behavior during a high-availability failover: the active device owns the floating address, and the same translated source address can move to the peer if the traffic group fails over. Therefore, for traffic leaving through the internal VLAN, BIG-IP first selects the internal floating self IP, **192.1.1.1**.

This selection is based on the VLAN used for the server-side egress path, rather than the virtual server address or the default route used for traffic directed externally. F5 documents the Auto Map selection behavior, including its use of self IP addresses and preference for floating addresses, in [K13349: SNAT automap source address selection](#). See also the F5 [Source Address Translation documentation](#).

QUESTION NO: 7

Refer to the exhibit.

A BIG-IP Administrator creates a new Virtual Server. The end user is unable to access the page. During troubleshooting, the administrator learns that the connection between the BIG-IP system and server is NOT set up correctly.

What should the administrator do to solve this issue?

- A. Disable Address Translation
- B. Set Address Translation to Auto Map, configure a SNAT pool, and have pool members in the same subnet of the servers
- C. Set Address Translation to SNAT and configure a specific translation address
- D. Set Address Translation to SNAT and have self-IP configured in the same subnet of servers

ANSWER: C

Explanation:

Set Address Translation to SNAT and configure a specific translation address. Source Network Address Translation ensures that server-side connections initiated by BIG-IP use a source IP address that is routable by the pool members for their return traffic. This is necessary when the servers do not have a route back to the original client network, which is a common cause of a virtual server accepting a client connection while the application response never returns successfully through BIG-IP.

Selecting SNAT and specifying a translation address gives the administrator deterministic control over the source address presented to the servers. The selected translation address must be reachable from the server network and must have a return path through BIG-IP, so return packets are correctly associated with the existing flow and forwarded to the client. This approach also avoids requiring changes to every pool member's routing table when client source networks are numerous or unknown. F5 documents source address translation as a virtual-server setting and supports use of a configured translation address or SNAT pool for server-side traffic. See the [BIG-IP virtual server tmsh reference](#) and F5's [Source Address Translation documentation](#).

QUESTION NO: 8

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 establish an FTP control-session login and perform the configured file check. The **File path** supplies the target file name or path used by the monitor after it connects to the FTP service. This gives the monitor a specific resource to request and use as the basis for a successful health-check result. The **username** and **password** provide the credentials sent during FTP authentication. They must identify an account that is permitted to log in to the FTP server and access the configured file path. Together, these settings allow BIG-IP to verify not merely that a TCP listener responds, but that the FTP service can authenticate a client and retrieve the intended file resource. The pool member's configured service address and port are ordinarily used as the monitor destination; alias settings are available when an administrator needs to override that normal destination behavior, rather than being required information for every FTP monitor. F5 documents the FTP monitor's `filename`, `username`, and `password` settings in the [tmsh FTP monitor reference](#). See also F5's [FTP monitor documentation](#) for monitor behavior and configuration details.

QUESTION NO: 9

A webserver is being overloaded with HTTPS traffic. To decrease the load on the server, the LTM Specialist and the Server Administrator decide to perform SSL offloading on the LTM device. The configuration of the virtual server is as follows:

```

ltn virtual example.com vs {
  destination 10.10.10.1:https
  ip-protocol tcp
  mask 255.255.255.255
  pool webserver_pool
  profiles {
    clientssl {
      context clientside
    }
    http {
    }
    serverssl {
      context serverside
    }
    tcp {
    }
  }
  source 0.0.0.0/0
  vs-index 12
}

```

Which change must be made to the configuration to perform SSL offloading?

- A. Remove the clientssl and http profiles
- B. Remove the clients profile
- C. Remove the clientssl and serverssl profiles
- D. Remove the serverssl profile

ANSWER: D

Explanation:

Remove the serverssl profile. SSL offloading means that the BIG-IP system terminates the client-side TLS/SSL connection, decrypts the HTTP request, and forwards that request to the pool member without establishing a new SSL session on the server side. The clientssl profile remains attached to the virtual server because it supplies the certificate and TLS settings that allow BIG-IP to accept and decrypt HTTPS traffic from clients. The HTTP profile can also remain, allowing BIG-IP to inspect and process the decrypted HTTP protocol traffic.

With both a clientssl profile and a serverssl profile assigned, BIG-IP performs SSL bridging: it decrypts the client connection and then creates a separate encrypted SSL/TLS connection to the backend server. Although this supports encrypted traffic on both legs, the backend still performs TLS processing. Removing the serverssl profile makes the server-side connection clear-text HTTP, shifting the TLS handshake and encryption/decryption workload away from the web server and onto the LTM device. The pool members must therefore be configured to accept HTTP on the applicable service port. F5 documents client-side and server-side SSL profile behavior in its [SSL Profiles Reference](#) and SSL offload configuration concepts in the [BIG-IP LTM implementation guide](#).

QUESTION NO: 10

An LTM Specialist wants to allow access to the Always On Management (AOM) from the network.

Which two methods should the LTM Specialist use to configure the AOM interface? (Choose two.)

- A. Configure the AOM IP from the front panel buttons and LCD.

- B. Choose the network configurator in the AOM menu on the serial port.
- C. Configure the AOM network address in the GUI under System > Platform.
- D. Log in to the Host via ssh, "ssh aom", and modify the network configuration file.

ANSWER: B D

Explanation:

The appropriate methods are to choose the network configurator in the AOM menu on the serial port and to log in to the host through SSH, use `ssh aom`, and modify the network configuration file. AOM is an independent, out-of-band management subsystem that remains available even when the primary BIG-IP host is unavailable. It therefore has its own network settings, including an address, netmask, gateway, and DNS-related settings where applicable. Direct serial-console access provides access to the AOM menu and its network configuration utility, making it a reliable method during initial setup or recovery. When the BIG-IP host is reachable, an administrator can also connect from the host to the AOM subsystem with `ssh aom` and configure its network parameters through the AOM environment. After the AOM network interface is configured and connected, administrators can remotely reach the AOM for platform-level monitoring, console access, and recovery operations. This separation of the AOM management path from normal BIG-IP management is the key reason these configuration methods are used. See the F5 [BIG-IP System Platform Guide](#) and the [F5 Support knowledge base](#) for platform and management guidance.

QUESTION NO: 11

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:

configure a VLAN and primary mirroring address for mirror traffic is required because BIG-IP uses a configured mirroring address to send connection-state and persistence-state updates from the active device to its peer. Assigning that address on an appropriate VLAN provides the dedicated, reachable path needed for reliable state replication between the high-availability devices. The peer can then possess current flow information before a failover occurs.

enable Mirroring for a virtual server and persistence profile is also required. Virtual-server connection mirroring replicates the state of eligible active TCP connections, allowing the newly active device to continue handling established client/server flows after takeover. Enabling persistence mirroring on the Universal Persistence profile replicates persistence records so that subsequent requests continue to resolve to the same selected pool member. Together, connection mirroring and persistence mirroring address both already-established sessions and persistence affinity during failover, which is especially important where source translation is not in use and disruption must be minimized. F5 documents the mirroring-address configuration and connection-mirroring behavior in the [BIG-IP LTM High Availability documentation](#) and persistence mirroring in the [BIG-IP LTM Profiles Reference](#).

QUESTION NO: 12

An LTM Specialist regularly provides analytics reports that show that traffic generated by different subnets within the organization. The LTM Specialist needs show the associate department names next the IP addresses in the reports.

Which step should the LTM Specialist take to meet this requirement?

- A. use an iRule to change the output of the report
- B. export the report and add the department names manually
- C. create VLANs for each subnet and set the name accordingly
- D. define active subnets and assign a name to specific subnets

ANSWER: D

Explanation:

Define active subnets and assign a name to specific subnets. BIG-IP Analytics supports associating descriptive names with configured client subnets so that Analytics reports can display a meaningful subnet label alongside the relevant client IP address information. The specialist can use each department name as the subnet name and define the corresponding CIDR range for that department. This creates a reusable, centrally managed mapping rather than requiring a report-by-report modification. When Analytics processes client traffic, addresses that fall within a configured subnet are identified using the assigned subnet name, making reports easier for business stakeholders to interpret. For example, a range used by Finance can be labeled "Finance," while a range used by Engineering can be labeled "Engineering." This feature is particularly useful when report consumers need organizational context but may not recognize individual IP ranges. The Analytics implementation documentation describes configuring Analytics data collection and reporting behavior, including client-address-related reporting configuration. See the [BIG-IP Analytics implementation documentation](#) and the [F5 Analytics documentation](#).

QUESTION NO: 13

```
set payload {CACHE :: payload}
}
```

Which two profiles should be used on the virtual server? (Choose two.)

- A. http-transparent
- B. http compression
- C. http
- D. webacceleration
- E. stream

ANSWER: C D

Explanation:

The `http` and `webacceleration` profiles are required for a virtual server using the `CACHE :: payload` iRule command. `CACHE :: payload` operates on the object body maintained by the BIG-IP web acceleration cache. The web acceleration profile enables the caching feature and supplies the RAM Cache behavior that allows BIG-IP to store HTTP objects, evaluate their cacheability, and serve cached responses. It is therefore the profile that makes cached payload data available to the iRule.

An HTTP profile is also needed because web acceleration processing is HTTP-aware. BIG-IP must parse the client request and server response as HTTP transactions, including headers, methods, response status, and object metadata, before it can apply cache processing and expose cached content to iRules. With both profiles associated with the virtual server, the iRule can access the cache payload in the appropriate cache-processing context and use it for the intended response-handling logic. F5 documents `CACHE :: payload` as an iRules cache command and describes web acceleration as the BIG-IP profile

used for HTTP content caching. See the [F5 CACHE::payload iRules reference](#) and the [F5 web acceleration profile 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 web 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:

Remove HTTP compression profile and **Remove web acceleration profile.** are the appropriate changes because both features add application-layer processing work to BIG-IP for every eligible HTTP response. HTTP compression requires BIG-IP to inspect response content and dynamically compress it when the client and policy permit, which consumes processing resources as traffic volume rises. Web acceleration similarly performs HTTP optimizations and caching-related processing intended to improve client response times, but it can add load to the device under a large request rate. Removing these optional optimization services reduces the LTM processing burden while retaining the functionality required by the application. In particular, the HTTP profile remains available for HTTP-aware processing, and SSL termination can remain in place so BIG-IP can read and act on the HTTP cookie used for persistence. Cookie persistence therefore continues to direct a returning client to the appropriate pool member, while the device no longer spends resources on compression and web-acceleration functions. F5 documents HTTP compression as an HTTP profile feature and describes web acceleration as a separate optimization capability that can be associated with virtual servers. See [F5 HTTP Compression Profile documentation](#) and [F5 BIG-IP WebAccelerator documentation](#).

QUESTION NO: 15

-- Exhibit --

-- 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 designed as an allowlist for HTTP response headers: BIG-IP retains the headers specified in the field and removes response headers that are not explicitly allowed. Therefore, specifying `Set-Cookie` and `Location` achieves the stated requirement to send only those two server response headers to the client. This is preferable when the goal is to preserve a small, known set of required headers while suppressing every other header, because the policy remains effective even if the origin server begins sending additional headers later.

In the HTTP profile, response-header controls are applied to traffic flowing from the server side through BIG-IP toward the client. The header names should be entered according to the profile field's supported format, typically as the configured list of allowed headers. `Set-Cookie` is required to deliver application session or preference cookies, while `Location` is required when the application returns redirect responses. F5 documents HTTP profile settings, including response-header manipulation controls, in its BIG-IP LTM documentation: [F5 CloudDocs: HTTP profile settings](#) and [F5 TechDocs: HTTP profile reference](#).

QUESTION NO: 16

Which three HTTP headers allow an application server to determine the client's language compatibility, browser, operating system type, and compression compatibility? (Choose three.)

- A. Accept
- B. Accept-Encoding
- C. Accept-Language
- D. Host
- E. User-Agent

ANSWER: B C E

Explanation:

Accept-Encoding, **Accept-Language**, and **User-Agent** provide the requested client-capability and client-environment information. **Accept-Language** communicates the user agent's preferred natural languages and associated preference weights, allowing an application to select an appropriate localized representation. **Accept-Encoding** identifies the content-coding algorithms that the client can decode, such as gzip, enabling the server to choose a compatible compressed response. **User-Agent** contains identifying information about the requesting user agent and commonly includes browser and operating-system details. Although applications should not rely on User-Agent parsing as the sole basis for feature decisions because values can be spoofed or reduced for privacy, it is the header conventionally used to identify the browser and platform characteristics described in the question. Together, these three request headers enable server-side selection based on language preference, supported response compression, and the browser/operating-system context. Header definitions and their intended negotiation behavior are documented by MDN for [Accept-Language](#) and [Accept-Encoding](#).

QUESTION NO: 17

A BIG-IP Administrator needs to make sure that the automatic update check feature works properly.

What must the administrator configure on the BIG-IP system?

- A. Update Check Schedule
- B. NTP servers
- C. DNS name servers
- D. SMTP servers

ANSWER: C

Explanation:

DNS name servers must be configured because the BIG-IP automatic update check contacts F5 update infrastructure using host names. Before BIG-IP can initiate that outbound update query, it must be able to resolve the relevant F5 service names to IP addresses. Configuring one or more reachable DNS resolvers under the system DNS settings provides that name-resolution capability and is therefore a required foundational setting for a reliable automatic update check. The configured resolvers should be reachable through the appropriate management or data-plane route, according to the deployment's networking design. Administrators can configure these values in the BIG-IP Configuration utility or with TMOS commands such as `tmsh modify sys dns name-servers { ... }`. The BIG-IP TMOS reference describes the `sys dns` object and its `name-servers` setting at [F5 BIG-IP TMOS sys dns reference](#). F5's DNS configuration guidance is also available in the [BIG-IP System Essentials documentation](#). Once DNS resolution is available, the automatic update-check service can resolve its update destination and perform its scheduled check.

QUESTION NO: 18

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. Server latency is the HTTP transaction timing metric that records the time associated with the server-side portion of a request: after BIG-IP forwards the request toward the pool member and until it receives the server response. It therefore provides the relevant measurement for delay between the BIG-IP system and the internal web server, rather than an end-to-end client experience measurement. Client IP collection adds the client address as an Analytics reporting dimension. With it enabled, the administrator can filter or group the collected server-latency data for the particular client in question, instead of seeing only aggregate latency across all clients. Together, these settings allow the specialist to identify that client's requests and evaluate the internal server-response delay associated with them. F5 documents HTTP Analytics data collection and reporting configuration in the [BIG-IP Analytics implementation documentation](#); related BIG-IP Analytics guidance is also available through [F5 BIG-IP Analytics documentation](#).

QUESTION NO: 19

A BIG-IP Administrator recently deployed an application. Users are experiencing slow performance with the application on some remote networks.

Which two modifications can the BIG-IP Administrator make to address this issue? (Choose two)

- A. Apply dest_addr profile to the Virtual Server
- B. Apply f5-tcp-wan profile to the Virtual Server
- C. Apply f5-tcp-lan profile to the Virtual Server
- D. Apply source_addr profile to the Virtual Server

E. Apply fasti_4 profile to the Virtual Server

ANSWER: B C

Explanation:

Applying **f5-tcp-wan** and **f5-tcp-lan** profiles tunes BIG-IP TCP behavior for the distinct network characteristics on each side of a full-proxy connection. Remote users commonly connect across higher-latency, lower-bandwidth, and potentially lossy WAN paths. A WAN-oriented TCP profile uses TCP settings intended to maintain effective throughput and recovery over those conditions, including more suitable buffering and congestion-control-related behavior than a profile intended solely for short, fast local links.

A LAN-oriented TCP profile complements that optimization for the connection between BIG-IP and application servers, where latency is normally low and bandwidth is relatively high. Using the appropriate client-side and server-side TCP profiles lets BIG-IP optimize each independent TCP flow according to its actual path, rather than forcing one set of TCP assumptions on both the remote-client and local-server networks. This is a core advantage of BIG-IP full proxying and is directly relevant when performance degradation is observed only from remote networks. F5 documents TCP profiles and their client-side/server-side contexts in the [BIG-IP TCP profile API reference](#); TCP profile configuration concepts are also covered in the [BIG-IP TCP Profile reference](#).

QUESTION NO: 20

A vulnerability scanner sends test requests to an ASM-protected application. One attack signature produces expected alerts for the scanner, but it also blocks a required scanner check. The signature must remain enforced for public client traffic and for all other application URLs.

Which configuration provides the most targeted solution?

- A. Create a signature exception limited to the affected scanner request context.
- B. Disable the attack signature for the entire security policy.
- C. Change the security policy to Transparent enforcement mode.
- D. Add the scanner source address as trusted for all policy checks.

ANSWER: A


Explanation:

A signature exception scoped to the scanner request context provides the narrowest solution. The exception can be limited to the applicable URL, parameter, or other supported context so the required scanner request is not blocked while the signature remains enabled for the rest of the application.

Disabling the signature in the entire policy removes protection for every URL. Placing the entire policy in Transparent mode prevents blocking for all violations, and adding the scanner address as a trusted source can bypass more protection than is required for one known signature match.

QUESTION NO: 21

Refer to the Exhibit.



An LTM Specialist notices that two members in a pool are overloaded. To relieve the existing members a fourth member (10.128.20.14) is brought up.

How many member will receive and process new connections?

- A. 4
- B. 3
- C. 2
- D. 1

ANSWER: C

Explanation:

2 is correct. The pool configuration shown in the exhibit uses priority group activation, which determines which pool members are eligible for new connections according to their assigned priority group and the configured minimum number of active members. BIG-IP first uses the highest-priority group that can satisfy the activation requirement. Once that requirement is met, members in lower-priority groups remain available as standby capacity but are not selected for new connections. Therefore, bringing up member 10.128.20.14 does not automatically make every available pool member eligible to receive traffic. With the configured priority-group activation threshold satisfied by two members in the active priority group, exactly two members receive and process new client connections. This behavior lets administrators reserve lower-priority members for failover or additional capacity when the preferred group no longer has enough healthy members. The BIG-IP LTM pool API documents the priority-group and minimum-active-members settings that control this selection behavior: [F5 BIG-IP LTM pool API reference](#). F5 also describes priority group activation as a mechanism for controlling which pool members are active for load balancing: [F5 priority group activation documentation](#).

QUESTION NO: 22

An LTM device has a virtual server mapped to www5f.com with a pool assigned. The objects are defined

as follows:

Virtual server. Destination 192.168.245.100.443 netmask 255.255.255.0

Persistence: Source address persistence netmask 255.0.0.0

SNAT:AutoMap

Profiles: HnP/TCP

How should the BIG-IP Administrator modify the persistence profile so that each unique IP address creates a persistence record?

- A. netmask 0.0.0.0
- B. netmask 255.255.255.255
- C. netmask 255.255.0.0
- D. netmask 255.256.255.0

ANSWER: B

Explanation:

netmask 255.255.255.255 is correct because a source-address persistence profile uses its configured netmask to determine how much of the client source address is used as the persistence key. A /32 mask, represented as 255.255.255.255, retains every bit of an IPv4 client address. Consequently, each distinct client IPv4 address produces its own persistence record and is consistently associated with the selected pool member for the profile's configured timeout.

The current 255.0.0.0 mask is a /8 mask, so persistence is based only on the first octet of the client address rather than the complete address. Changing the source-address persistence mask to 255.255.255.255 provides host-specific persistence. SNAT Auto Map does not alter this requirement: BIG-IP evaluates the original client source address for source-address persistence before source address translation is applied toward the pool member. This configuration is appropriate when the objective is individual client affinity rather than grouping clients by a network prefix.

F5 documents source-address persistence and its mask setting in the [tmsh persistence source-address reference](#). General persistence-profile behavior is also described in the [BIG-IP LTM persistence profiles 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 gateway_icmp monitor to the pool
- E. enable manual-resume on the same monitor

ANSWER: B D

Explanation:

Adding a receive string to the game monitor makes the UDP health check response-based rather than relying solely on successful transmission of a UDP probe. UDP is connectionless, so a probe can leave BIG-IP successfully even when the server-side switch is unavailable. With a receive string configured, the monitor must receive the expected application response within its configured timeout. When the access switch is down, the response cannot return, the monitor times out, and BIG-IP marks the member unavailable.

Also assigning a gateway_icmp monitor to the pool adds an ICMP reachability check for the relevant gateway path. A server-side network failure that prevents traffic from reaching the member or its configured gateway causes the ICMP monitor to fail after its timeout. When both monitors are associated under the applicable monitor rule, BIG-IP evaluates their availability results and can remove the member from service, allowing the pool to become unavailable and the configured disaster-recovery behavior to occur. The timeout must remain longer than the interval so BIG-IP has sufficient time to evaluate a missed response. F5 documents UDP monitor receive-string behavior in the [UDP monitor API reference](#) and gateway reachability monitoring in the [gateway_icmp monitor API reference](#).

QUESTION NO: 24

Which two alerting capabilities can be enabled from within an application visibility reporting (AVR) analytics profile? (Choose two.)

- A. sFlow
- B. SNMP
- C. e-mail
- D. LCD panel alert
- E. high speed logging (HSL)

ANSWER: B C

Explanation:

Within an AVR analytics profile, alert rules can be configured to notify operations staff when collected application-traffic metrics cross defined thresholds. The supported notification actions include **SNMP**, which sends an SNMP trap to a configured SNMP destination, and **e-mail**, which sends an alert message through the BIG-IP system's configured SMTP facilities. These actions allow analytics data such as traffic volume, server-side response time, client-side response time, transactions per second, or request-related conditions to produce actionable operational notifications. The alert is defined in the Analytics profile and becomes active when that profile is associated with the relevant virtual server, making notifications specific to the monitored application flow. F5 documents Analytics alerts as having selectable alert actions, including sending an SNMP trap and sending an email. The email action requires appropriate SMTP configuration on the BIG-IP system, while the SNMP action requires SNMP/trap destination configuration so the receiving monitoring platform can process the notification. See F5's [Analytics implementation documentation](#) and the [F5 SNMP reference](#) for related monitoring and notification configuration information.

QUESTION NO: 25

Which command should an LTM Specialist use on the command line interface to show the health of RAID array hard drives?

- A. tmsh show /sys raid disk
- B. tmsh show /ltm raid disk
- C. tmsh show /sys raid status
- D. tmsh show /ltm disk status

ANSWER: A

Explanation:

`tmsh show /sys raid disk` is the appropriate command because RAID disk objects are managed and reported within the BIG-IP system (`/sys`) hierarchy. The command displays per-disk RAID information, allowing an LTM Specialist to inspect the operational condition of the physical drives that participate in the appliance RAID configuration. This is the relevant level of detail when the goal is to assess hard-drive health rather than merely view a high-level array summary. Disk-level output is useful during hardware troubleshooting because it identifies the individual drive state and helps determine whether a drive is online, unavailable, rebuilding, or otherwise requires attention. The command must be run from the BIG-IP command line with sufficient privileges to access system-level status information. BIG-IP separates traffic-management objects, such as virtual servers and pools, from platform and hardware objects; RAID administration is therefore exposed through the system namespace. F5's TMOS Shell reference documents the `show` command structure and system-object hierarchy, while the BIG-IP system documentation describes monitoring platform hardware and storage status. See the [F5 TMSH Reference](#) and [F5 BIG-IP System Essentials documentation](#).

QUESTION NO: 26

An LTM Specialist needs to create an iRule that creates persistence records based on a JSESSIONID cookie. If a persistence record already exists, then the iRule must persist the client connection according to the existing record.

Which persistence profile enables the iRule to meet these requirements?

- A. Universal
- B. SSL
- C. Destination Address Affinity
- D. Cookie
- E. Source Address Affinity

ANSWER: A

Explanation:

Universal is correct because a Universal persistence profile is specifically designed to work with an iRule that supplies a custom persistence key. The iRule can extract the value of the JSESSIONID cookie and use the `persist uie` command to create or look up a Universal Persistence (UIE) record using that value. When the same JSESSIONID value is presented on a later request, BIG-IP finds the existing UIE record and directs the connection to the pool member associated with that record. This supports application-defined session affinity while allowing the iRule to control precisely how the cookie is parsed, normalized, and used as the persistence key. A Universal persistence profile must be associated with the virtual server so BIG-IP has the UIE persistence method and timeout settings available when the iRule invokes persistence. The iRules `persist` command documentation describes the `uie` mode for using a user-defined persistence key, and F5's Universal persistence guidance identifies it as the profile type used for custom iRule-based persistence logic. See the [F5 iRules persist command reference](#) and [F5 Universal persistence documentation](#).

QUESTION NO: 27

A legacy search page accepts a query parameter named `q`. A particular enabled attack signature generates false positives only when the signature matches the `q` parameter on `/legacy/search`. The same signature has blocked confirmed malicious traffic on other URLs.

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

- A. Create a signature exception for the `q` parameter on `/legacy/search`.
- B. Keep the signature enabled and enforced for the remaining policy scope.
- C. Disable the signature for the entire security policy.

D. Change the entire policy to Transparent enforcement mode.

E. Allow all meta characters for every parameter in the policy.

ANSWER: A B

Explanation:

The administrator should create a signature exception limited to the affected URL and parameter context. This prevents the known false positive only for `q` on `/legacy/search`.

The signature should remain enabled and enforced for the rest of the policy because it has detected confirmed attacks elsewhere. A narrowly scoped exception preserves that protection while accommodating the documented application behavior.

Disabling the signature globally, placing the entire policy in Transparent mode, or allowing all meta characters for every parameter would create a much broader reduction in protection than the scenario requires.

QUESTION NO: 28

An LTM device is deployed in a one-armed topology. The virtual server, clients, and web servers are connected on the LTM device internal VLAN. A client tries to connect to the virtual server and is unable to establish a connection. A packet capture from the LTM device internal VLAN shows that the HTTP request is being forwarded to the web server.

From which two additional locations should protocol analyzer data be collected? (Choose two.)

- A. network interface of web server
- B. network interface of client machine
- C. internal VLAN interface of LTM device
- D. external VLAN interface of LTM device
- E. any network interface of the Internet firewall

ANSWER: A B

Explanation:

The appropriate additional capture points are the *network interface of web server* and the *network interface of client machine*. Seeing the HTTP request leave BIG-IP on the internal VLAN confirms that BIG-IP selected a pool member and forwarded the client-side request toward that server. A simultaneous capture at the web server NIC establishes whether the request actually arrives at the server and whether the server emits TCP acknowledgments and an HTTP response. This distinguishes an issue on the server host or its local network attachment from a return-path issue.

A capture at the client NIC completes the end-to-end view. It verifies the client-to-virtual-server TCP exchange, confirms what the client actually transmitted, and determines whether any response packets reach the client. In a one-armed design, client and server traffic uses the same connected VLAN, so checking both endpoints alongside the existing BIG-IP VLAN capture is especially useful for detecting asymmetric return traffic, direct server-to-client responses, or packets that are present on the wire but not received or processed by an endpoint. F5 recommends packet captures at relevant traffic points when tracing connection flow; see [F5 KB K411: Overview of packet tracing with tcpdump](#) and [F5 BIG-IP Local Traffic Management documentation](#).

QUESTION NO: 29

-- Exhibit --

-- Exhibit --

Refer to the exhibits.

Which two servers are missing two frequently used URLs? (Choose two.)

- A. 172.16.20.1 /text.one /text.txt
- B. 172.16.20.2 /text.one /text.txt
- C. 172.16.20.1 /text.txt /browserspecific.html
- D. 172.16.20.2 /text.one /browserspecific.html
- E. 172.16.20.3 /text.one /browserspecific.html

ANSWER: B E

Explanation:

The servers identified by comparing the frequently used URL list with the URLs shown for each server are **172.16.20.2 /text.one /text.txt** and **172.16.20.3 /text.one /browserspecific.html**. For 172.16.20.2, both `/text.one` and `/text.txt` appear in the frequently used URL information but are absent from that server's recorded URL set. For 172.16.20.3, the absent frequently used URLs are `/text.one` and `/browserspecific.html`. This type of comparison is relevant when reviewing per-server application behavior: an ASM policy can learn and enforce expected URLs, while differences in URL availability between pool members can indicate inconsistent application deployment or a server-specific content issue. The appropriate operational response is to validate that the expected application content and URI handling are deployed consistently across the affected pool members before relying on a learned policy to enforce the application's URL inventory. F5 documentation describes URL-related policy entities and enforcement as part of application-security policy configuration and management. See the [F5 BIG-IP ASM implementation documentation](#) and the [F5 BIG-IP ASM technical documentation](#).

QUESTION NO: 30

A policy blocks requests with the violation *Illegal parameter data type*. The application team states that the affected parameter, `accountNumber`, is valid but now contains alphanumeric account identifiers rather than the numeric values used when the policy was created.

Which two policy settings should the ASM Specialist review? (Choose two.)

- A. The configured data type for `accountNumber`
- B. The allowed meta characters for `accountNumber`
- C. The pool health monitor assigned to the application servers
- D. The TCP profile assigned to the virtual server
- E. The virtual server destination address

ANSWER: A B

Explanation:

The configured data type for `accountNumber` determines the values ASM considers valid. If the parameter is defined as numeric, an alphanumeric identifier can correctly trigger an *Illegal parameter data type* violation. The parameter data type should be reviewed and changed only if the new application behavior is valid and intended.

Allowed meta characters should also be reviewed. Even after selecting an appropriate data type, the parameter may require explicitly permitted characters, such as letters or a hyphen, that are part of the new identifier format.

The pool health monitor and TCP profile do not parse HTTP parameter values. Changing the virtual server destination would not alter ASM parameter validation.

QUESTION NO: 31

Which two subsystems could the LTM Specialist utilize to access an LTM device with lost management interface connectivity? (Choose two.)

- A. AOM
- B. ILO
- C. SCCP
- D. ALOM

ANSWER: A C


Explanation:

AOM and **SCCP** provide out-of-band paths for administering applicable BIG-IP hardware when the normal management interface cannot be reached. The Always-On Management subsystem is an independent service processor that remains available separately from the BIG-IP host operating environment. An administrator can use it for recovery-oriented access, including console access and, on supported platforms and configurations, dedicated management connectivity. This permits investigation and remediation even when the primary management plane is unavailable.

The Switch Card Control Processor is the chassis-management control processor used with applicable VIPRION systems. It provides an independent management and console path to the chassis infrastructure and supports access needed to diagnose or recover a blade or chassis when ordinary device management connectivity has failed. These subsystems are therefore valuable for restoring administrative control without relying on the unavailable BIG-IP management interface. F5 maintains platform-specific hardware and out-of-band management procedures in its [Technical Documentation](#) and publishes hardware-platform guidance through the [F5 CloudDocs documentation portal](#).

QUESTION NO: 32

-- Exhibit --



-- Exhibit --

Refer to the exhibit.

A layer 2 nPath routing configuration has been deployed. A packet capture contains a client connection packet with the following properties:

Source IP:

Destination IP:

At which two locations could the packet capture have been taken? (Choose two.)

A. the network interface of web server

- B. the DMZ interface of the Internet firewall
- C. the internal interface of the Internet firewall
- D. the external VLAN interface of the LTM device

ANSWER: A C

Explanation:

In a Layer 2 nPath deployment, BIG-IP forwards the inbound client flow to the selected server without performing the normal full-proxy source and destination address translations. The server therefore receives traffic addressed to the virtual server address and sends its response using that virtual server address as the packet source, with the client's address as the destination. This return packet is emitted by the web server, so it can be captured directly on the network interface of web server.

nPath is specifically designed to permit a return path that bypasses BIG-IP. In the illustrated topology, the server's return traffic proceeds toward the client through the routing and security infrastructure on the server side. Consequently, the same packet, retaining the virtual server source IP and Client A destination IP, can be observed on the internal interface of the Internet firewall. The packet's IP addresses are significant because they identify it as server-to-client return traffic for a connection accepted through the virtual server, rather than the original client-to-virtual-server request.

F5 describes nPath as a forwarding method in which BIG-IP directs client traffic to a server while server responses can use an alternate return path. See [F5 KB K10040: Overview of nPath routing](#) and the [BIG-IP Local Traffic Management documentation](#).

QUESTION NO: 33

An ASM Specialist uses manual learning for a new security policy. After reviewing a learning suggestion for a legitimate URL, the specialist approves the suggestion. The policy must begin enforcing the resulting URL configuration.

Which two actions are required? (Choose two.)

- A. Accept the approved learning suggestion.
- B. Change the virtual server pool.
- C. Restart the ASM service.
- D. Apply the security policy.
- E. Change the policy to Transparent enforcement mode.

ANSWER: A D

Explanation:

The approved learning suggestion must be accepted so that the learned URL becomes part of the policy configuration. The policy must then be applied for the accepted configuration changes to become active in ASM enforcement.

Changing the virtual server pool or restarting the ASM service does not commit a learning suggestion into the policy. Switching the policy to transparent mode would prevent blocking rather than activate the approved enforcement configuration.

QUESTION NO: 34

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 could cause this behavior because BIG-IP can monitor a VLAN for received traffic. When VLAN failsafe is enabled and the configured timeout expires without traffic being received on that VLAN, the device treats the condition as a failsafe event. On an active device in a redundant device group, the configured failsafe action can initiate failover. Therefore, an upstream switch outage that interrupts traffic on the monitored VLAN can directly coincide with and trigger the active unit's transition to standby.

a monitor on a pool in an HA group could also lead to failover. An HA group assigns a health score to each device using configured resources, including pools and their available members. Pool-member availability is determined by the pool's health monitors. If a switch outage makes monitored pool members unreachable, the active device's HA-group score can fall below the peer's score. BIG-IP can then fail over to the device with the better HA-group score. This mechanism lets failover account for application and network reachability, rather than relying solely on device heartbeat status.

F5 documents VLAN failsafe controls in the [tmsh net vlan reference](#) and HA-group scoring configuration in the [tmsh ltm ha-group reference](#).

QUESTION NO: 35

A user is having issues with connectivity to an HTTPS virtual server. The virtual server is on the LTM device's external vlan, and the pools associated with the virtual server are on the internal vlan. An LTM Specialist does a tcpdump on the external interface and notices that the host header is incomplete.

In which location should the LTM Specialist put a traffic analyzer to gather the most pertinent data?

- A. server
- B. external VLAN
- C. internal VLAN
- D. client machine

ANSWER: D

Explanation:

client machine is correct because the HTTP `Host` header is application-layer content carried within the HTTPS session. HTTPS protects the HTTP request with TLS while it traverses the network, so a packet capture taken from the external network path does not necessarily provide a complete, directly readable representation of the original HTTP request. Capturing or analyzing traffic at the client endpoint provides the most useful evidence of what the client application actually generated before TLS protection is applied, including the intended hostname and request construction. This makes it possible to distinguish an application-side request issue from a transport or BIG-IP processing issue and to correlate the client's request timing with the connection attempt. F5 documents that HTTPS virtual servers use SSL profiles to process encrypted client-side traffic, and that packet captures are a primary diagnostic tool whose usefulness depends on the capture point and available decrypted data. See F5's [SSL offload documentation](#) and [F5 tcpdump guidance](#).

QUESTION NO: 36

A BIG-IP Administrator needs to check the memory utilization on a BIG-IP system. Which two methods can the BIG-IP Administrator use? (Choose two.)

- A. Run the `tmsh show/sys memory` command
- B. Run the `tmsh show/sys traffic` command
- C. Go to Statistics > Module Statistics > Traffic Summary in the configuration utility
- D. Go to Statistics > Module Statistics > Memory in the configuration utility
- E. Go to System > Disk Management in the configuration utility

ANSWER: A D

Explanation:

Run the `tmsh show/sys memory` command is a direct command-line method for viewing BIG-IP memory statistics. From the Traffic Management Shell, the `show sys memory` command displays system memory information, including total memory, memory currently in use, and free memory. This makes it useful for an administrator connected through SSH or the serial console who needs to assess system resource consumption without using the Configuration utility. F5 documents the `sys memory` object and its statistics in the [tmsh sys reference](#).

Go to Statistics > Module Statistics > Memory in the configuration utility is the corresponding graphical method. The Memory statistics view presents memory-utilization data through the BIG-IP web interface, allowing an administrator to inspect memory consumption visually as part of routine system monitoring. This is appropriate when operational access is through the Configuration utility or when the administrator wants to review the platform's statistics in the GUI. F5's system monitoring documentation describes using BIG-IP statistics and monitoring views to observe system health and resource utilization: [BIG-IP System Monitoring and Management](#).

QUESTION NO: 37

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 every BIG-IP traffic group is a uniquely named configuration object. A traffic group identifies the set of floating configuration objects—such as virtual addresses and virtual servers—that can move together from one device to another within a device service cluster when failover occurs. When an administrator creates an additional traffic group, BIG-IP must have a distinct name by which it can store, reference, assign, and administer that group. The name also allows related objects to be explicitly associated with the intended traffic group rather than remaining in the default traffic group. In the BIG-IP TMOS command reference, the traffic-group object is created and managed by specifying its object name, confirming that a name is the fundamental required identifier for the new object. Optional operational settings can subsequently control behavior such as fallback and Layer 2 address handling, but the traffic group must first exist as a named object. See the F5 [tmsh cm traffic-group reference](#) and F5 [traffic group administration documentation](#).

QUESTION NO: 38

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.example.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. HTTPCompression
- D. HTTP
- E. TCP

ANSWER: D E

Explanation:

HTTP and **TCP** are the required profiles for this type of Layer 7 pool-selection configuration. The HTTP profile enables BIG-IP to parse the client request as HTTP and expose HTTP-specific request attributes, including the URI path. An administrator can then use an iRule, such as one triggered by `HTTP_REQUEST`, to evaluate `HTTP::path` and select the appropriate pool with the `pool` command. For example, requests whose path starts with `/app1` can be directed to the app1 pool, while requests beginning with `/app2` can be directed to the app2 pool. F5 documents that `HTTP::path` returns the path portion of an HTTP request URI and is used in HTTP iRule events: [F5 HTTP::path iRule command reference](#).

The TCP profile provides the transport-layer connection handling for the virtual server. HTTP runs over TCP, so BIG-IP needs a TCP profile to establish, proxy, and manage the client-side and server-side TCP flows carrying those HTTP requests. Together, TCP provides the Layer 4 transport foundation and HTTP provides the Layer 7 parsing needed to inspect the URL path and apply path-based pool selection. F5's iRules reference describes the TCP event context and commands available for TCP connections: [F5 TCP iRules reference](#).

QUESTION NO: 39

An LTM Specialist reports that an application is no longer reachable after it has been upgraded.

Nothing has been changed in the configuration on the LTM device.

The logs indicate that health monitors to all servers have failed as shown:

What should the LTM Specialist verify next?

- A. That the TCP hand shake with the servers is still completed using tcpdump
- B. That the custom receive string for the HTTP monitor has changed with the upgrade
- C. That the can still ping the servers from the BIG-IP device.
- D. That the firewall between the BIG-IP device and servers is still allowing HTTP

ANSWER: B

Explanation:

That the custom receive string for the HTTP monitor has changed with the upgrade is the appropriate next verification. A BIG-IP HTTP monitor does more than establish a TCP connection: it sends an HTTP request and evaluates the response against its configured receive string. A server is marked available only when the monitor receives a response matching that expected content. An application upgrade can change a page title, banner, redirect behavior, response body, URI handling,

or other HTTP response details without any change to the BIG-IP configuration. In that situation, transport connectivity can remain healthy while every pool member is marked down because the monitor's existing receive string no longer matches the upgraded application's response.

The specialist should inspect the actual HTTP response from a pool member and compare it with the monitor's configured send request and receive string. If the application's expected response has legitimately changed, update the monitor to use a stable and distinctive response value that confirms the application is functioning correctly. F5 documents that HTTP monitors use send and receive strings to determine availability and that a configured receive string must be found in the response. See the [F5 HTTP monitor reference](#) and the [F5 health monitor troubleshooting guidance](#).

QUESTION NO: 40

-- Exhibit --

-- Exhibit --

Refer to the exhibit.

An LTM Specialist is troubleshooting an issue with SSL and is receiving the error shown when connecting to the virtual server. When connecting directly to the pool member, clients do NOT receive this message, and the application functions correctly. The LTM Specialist exports the appropriate certificate and key from the pool member and imports them into the LTM device. The LTM Specialist then creates the Client SSL profile and associates it with the virtual server.

What is the issue?

- A. The SSL certificate and key have expired.
- B. The SSL certificate and key do NOT match.
- C. The client CANNOT verify the certification path.
- D. The common name on the SSL certificate does NOT match the hostname of the site.

ANSWER: C

Explanation:

The client CANNOT verify the certification path. Exporting and importing only the server certificate and its private key does not necessarily include the intermediate CA certificates that establish a complete trust chain from the server certificate to a trusted root CA. The pool member can work when accessed directly because it may have been configured to present the required intermediate certificates, while the BIG-IP Client SSL profile may present only the imported leaf certificate. When the virtual server terminates the TLS connection, BIG-IP becomes responsible for sending the appropriate certificate chain to connecting clients. The Client SSL profile should therefore be configured with a certificate chain file that includes the relevant issuing intermediate CA certificate or certificates, in the proper order. The client can then build and validate the chain from the presented server certificate through the intermediate issuer to a root CA in its trust store. F5 documents that Client SSL profiles use a certificate, key, and optional certificate chain to authenticate the BIG-IP system to clients. See the [F5 BIG-IP SSL Administration documentation](#) and the [X.509 PKI certificate-path validation standard \(RFC 5280\)](#).

QUESTION NO: 41

A pool of four servers has been partially upgraded for two new servers with more memory and CPU capacity. The BIG-IP Administrator must change the load balance method to consider more connections for the two new servers. Which load balancing method considers pool member CPU and memory load?

- A. Round Robin
- B. Dynamic Ratio
- C. Ratio
- D. Least Connection

ANSWER: B

Explanation:

Dynamic Ratio is the appropriate load-balancing method because it uses performance information reported by pool members to calculate a changing, relative distribution ratio. With the required dynamic-ratio monitoring configured, BIG-IP can receive system-load data from each server, including processor and memory utilization, and use those values to direct a larger share of new connections to the upgraded servers when they have more available capacity. The ratios are not fixed: they are continually recalculated as the members' reported load changes. This makes Dynamic Ratio suitable when servers in the same pool have unequal hardware resources or when their available capacity fluctuates during normal operation. The load-balancing decision therefore reflects actual member capacity rather than simply rotating requests, using an administrator-defined static weight, or counting existing connections. F5 describes Dynamic Ratio as a method that adjusts traffic distribution using performance metrics obtained through a dynamic ratio monitor. See the F5 [BIG-IP LTM load-balancing methods documentation](#) and the F5 [Dynamic Ratio load-balancing overview](#).

QUESTION NO: 42

Six servers have a varying number of connections that change based on the user load.

Which load balancing method should an LTM Specialist apply to divided the web application traffic to the servers on the relative performance trend?

- A. Least Sessions
- B. Least Connections
- C. Predictive
- D. Ratio

ANSWER: C

Explanation:

Predictive is the appropriate load-balancing method because it dynamically evaluates the connection-load trend of each pool member and uses that trend to select the member expected to perform best for the next connection. Rather than treating the present connection count as the sole decision point, the BIG-IP system observes whether a member's load is increasing or decreasing over time. This makes the method particularly suitable when the six servers experience changing, uneven connection levels driven by user demand.

In BIG-IP LTM, Predictive load balancing is designed to estimate a pool member's relative performance from recent connection behavior. A member whose active connections are trending downward can be preferred over one with a similarly low current count whose connections are rapidly increasing. That forward-looking selection behavior aligns directly with the requirement to distribute web traffic according to a relative performance trend. The method is configured as a pool load-balancing mode, allowing the LTM to continually adapt new-connection placement as server load conditions change. F5 documents pool load-balancing methods and their behavior in the [BIG-IP Local Traffic Management documentation](#); the available pool load-balancing mode values are also listed in the [iControl REST pool API reference](#).

QUESTION NO: 43

An LTM device needs to be configured a virtual server. The application requires SSL encryption from the client to the server and an X-Forwarded-For added by the LTM device.

Which virtual server type should the LTM Specialist use?

- A. Forwarding (IP)
- B. Performance (HTTP)
- C. Standard

- D. Stateless
- E. Performance

ANSWER: C

Explanation:

Standard is the appropriate virtual server type because it provides full-proxy, Layer 7 processing, allowing BIG-IP LTM to apply the profiles required for this flow. A Client SSL profile can decrypt and manage the client-side TLS connection, while a Server SSL profile establishes and encrypts the separate TLS connection from BIG-IP to the pool member. This enables encryption on both network segments while still allowing the system to inspect and modify the HTTP request between the two SSL sessions.

An HTTP profile attached to a Standard virtual server can insert the `X-Forwarded-For` header. This header carries the originating client IP address to the backend application, rather than leaving the application to see only the BIG-IP self IP used for the server-side connection. In the HTTP profile, the *Insert XForwarded For* setting enables this behavior. The Standard virtual server is therefore suited to the combined requirement for SSL processing and HTTP-header manipulation. F5 documents Standard virtual servers as supporting application delivery and profile-based processing, including SSL and HTTP profiles. See [F5 BIG-IP LTM virtual server implementation documentation](#) and [F5 HTTP header documentation](#).

QUESTION NO: 44

Refer to the exhibit.

Destination Address/Mask: 192.168.162.80
Service Port: 443 HTTPS
Notify Status to Virtual Address: ☒
Availability: Available (Enabled) - The virtual server is available
Synccookie Status: Off
State:
Configuration: Basic
Protocol: f5-tcp-lan
Protocol Profile (Client):
Protocol Profile (Server):
HTTP Profile:
HTTP Proxy Connect Profile:
FTP Profile:
RTSP Profile:
SSL Profile (Client): /Common/clientsl
SSL Profile (Server): /Common/serverssl
Available for Client SSL: /Common/clientslinsecure-compatible, /Common/clientsl-secure, /Common/clientsl2, /Common/crypto-server-default-clientsl
Available for Server SSL: /Common/apm-default-serverssl, /Common/crypto-client-default-serverssl, /Common/pool-default-serverssl, /Common/serverssl

A BIG-IP Administrator needs to deploy an application on the BIG-IP system to perform SSL offload and

re-encrypt the traffic to pool members.

During testing, users are unable to connect to the application.

What must the BIG-IP Administrator do to resolve the issue?

- A. Remove the configured SSL Profile (Client)
- B. Configure Protocol Profile (Server) as splitsession-default-tcp
- C. Enable Forward Proxy in the SSL Profile (Client)
- D. Configure an SSL Profile (Server)

ANSWER: D

Explanation:

Configure an SSL Profile (Server) is required because the stated design calls for both client-side SSL termination and server-side SSL re-encryption. A Client SSL profile lets the BIG-IP system accept the client TLS connection, present the virtual server certificate, and decrypt the incoming HTTPS request. After that termination occurs, BIG-IP needs a separate Server SSL profile to initiate a new TLS session from the virtual server to each selected pool member. Attaching the Server SSL profile therefore supplies the SSL/TLS settings used on the server-side leg, including the TLS versions, ciphers, certificate validation behavior, and—when needed—the server name indication sent to the pool member. This creates the intended encrypted path in both directions: client-to-BIG-IP and BIG-IP-to-pool-member. Without a Server SSL profile, the traffic sent to a pool member would not be re-encrypted as required; an HTTPS-only pool member cannot successfully process that clear-text server-side connection. F5 distinguishes Client SSL profiles, which process inbound client connections, from Server SSL profiles, which establish SSL connections to backend servers. See F5's [SSL profile documentation](#) and [SSL profiles overview](#).

QUESTION NO: 45

Users are unable to reach an application. The BIG-IP Administrator checks the Configuration Utility and observes that the Virtual Server has a red diamond in front of the status. What is causing this issue?

- A. All pool members are down.
- B. The Virtual Server is receiving HTTPS traffic over HTTP virtual.
- C. The Virtual Server is disabled.
- D. All pool members have been disabled.

ANSWER: A

Explanation:

All pool members are down. In the BIG-IP Configuration utility, a red diamond represents an offline object state. For a virtual server that forwards traffic to a pool, this commonly occurs when no eligible pool members remain available to receive connections. BIG-IP determines member availability through the configured health monitors and the members' operational state, then rolls that state up to the pool and the virtual server. When every member is marked down, the pool has no viable destination for client flows, so the dependent virtual server is shown as offline and users cannot reach the application through it. The administrator should review the pool-member status and monitor details, including the monitor's probe results, network reachability from BIG-IP to the servers, application listener status, and any recent server or routing changes. Restoring successful health-monitor responses for at least one enabled pool member returns an eligible target to the pool and allows the virtual server to resume serving traffic. F5 documents virtual-server and pool health monitoring behavior in the [BIG-IP Local Traffic Manager Monitors Reference](#) and its [Local Traffic Manager Basics documentation](#).

QUESTION NO: 46

These log entries can have different root causes:

Jun 28 05:01:21 LTM_A notice mcpd[27545]: 0107143a:5: CMI reconnect timer: enabled

Jun 28 05:01:21 LTM_A notice mcpd[27545]: 01071431:5: Attempting to connect to CMI peer 1.1.1.2 port 6699

Jun 28 05:01:21 LTM_A notice mcpd[27545]: 01071432:5: CMI peer connection established to 1.1.1.2 port 6699

Jun 28 05:01:26 LTM_A notice mcpd[27545]: 0107143a:5: CMI reconnect timer: disabled, all peers are connected

Which two commands should be used to obtain additional information on these entries? (Choose two.)

- A. `tmsh show /sys mcpd`
- B. `bigstart status mcpd`
- C. `tmsh modify /sys db log.mcpd.level value debug`
- D. `tmsh modify /sys db log.cmi.level value debug`

ANSWER: B C

Explanation:

`bigstart status mcpd` is appropriate because the Connection and Mirroring Infrastructure messages are emitted by the `mcpd` process. Confirming the daemon's current service state helps determine whether the peer reconnection coincided with an `mcpd` availability or process-state event. The command reports whether that critical BIG-IP service is running and is a useful first collection step when reviewing intermittent peer connectivity events.

`tmsh modify /sys db log.mcpd.level value debug` is also appropriate because it increases the logging verbosity for `mcpd`, the component reporting the reconnect timer and peer connection messages. Debug-level logging can provide the additional connection, retry, and state-transition detail needed to correlate the observed CMI reconnection with conditions affecting device communication. This setting should normally be enabled only while reproducing or investigating the issue, then returned to the normal logging level to avoid excessive log volume.

F5 documents system database variables through the TMSH `sys db` reference and documents BIG-IP service management commands in its operational guidance. See the [F5 TMSH sys db reference](#) and the [F5 BIG-IP service status guidance](#).

QUESTION NO: 47

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` while having the server team replace that content with `SERVER DOWN` provides a simple application-controlled maintenance mechanism. An HTTP monitor declares a pool member available only when its expected receive string is present in the monitor response. When the server administrator changes the contents returned by `/test.html`, the expected value is no longer returned and BIG-IP marks the member unavailable. This removes the member from new load-balancing selections without requiring an administrator to log in to BIG-IP.

Configuring the receive disable string as `SERVER DOWN` is the more explicit maintenance-state design. When the monitor receives this configured disable string, BIG-IP places the member into a disabled state rather than treating the condition solely as a generic monitor failure. The server team can therefore control pool-member availability by serving `SERVER DOWN` from the monitored URI, while restoring normal service by returning the normal successful monitor content. F5 documents the HTTP monitor receive string and receive disable string as response values that determine the monitored object's availability or disabled status. See [F5 HTTP monitor reference](#) and [F5 tmsh HTTP monitor documentation](#).

QUESTION NO: 48

Refer to the exhibit.

A pool member fails the monitor checks for about 30 minutes and then starts passing the monitor checks. New traffic is Not being sent to the pool member.

What is the likely reason for this problem?

- A. The pool member is disabled
- B. Monitor Type is TCP Half Open
- C. Manual resume is enabled
- D. Time Until Up is zero

ANSWER: C

Explanation:

Manual resume is enabled is correct. The Manual Resume setting is intended for a pool member that should not automatically return to service merely because its health monitor begins succeeding again. When a monitor marks that member down, BIG-IP requires an administrator to explicitly resume or enable the member after the underlying issue has been investigated and corrected. Consequently, the member can be successfully responding to its assigned monitor while still being withheld from new load-balanced connections. This behavior is particularly useful where an application may become network-reachable before it is operationally ready to receive production traffic, or where administrators require a deliberate validation step following an outage. The approximately 30-minute failed-monitor period establishes that the member had transitioned out of service; subsequent successful checks alone do not clear the manual-resume requirement. An administrator can return the member to service through the BIG-IP Configuration utility or TMSH after confirming readiness. F5 documents this pool-member behavior through the [tmsh ltm pool reference](#), and its health-monitor operation is covered in the [BIG-IP Local Traffic Manager monitors reference](#).

QUESTION NO: 49 - (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 virtual server type. It can use an HTTP profile and Layer 7 processing, allowing traffic to be inspected and load balanced based on the HTTP URI (for example, through an iRule, Local Traffic Policy, or other L7 logic).

Forwarding (Layer 2), **Stateless**, and **Performance (Layer 4)** do not provide the required full HTTP/URI inspection capability.

QUESTION NO: 50

A device on the network is configured with the same IP address as the management address of the active LTM device, causing the management GUI to be inaccessible.

Which two methods should the LTM Specialist use to access the LTM device in order to change the management IP address? (Choose two.)

- A. Connect via ssh to the AOM IP address.
- B. Connect via ssh to the management address.
- C. Connect to the LTM device via serial connection.
- D. Connect a monitor and keyboard to the LTM device.
- E. Connect via ssh to the standby unit and connect via ssh across the serial link between the devices.

ANSWER: A C

Explanation:

Connect via ssh to the AOM IP address. and **Connect to the LTM device via serial connection.** are the appropriate out-of-band access methods when the BIG-IP management IP address conflicts with another host. An address conflict makes traffic directed to the configured management address unreliable, so administrative access through that address cannot be depended upon to reach the intended active device. The Always-On Management (AOM) interface is an independent service-processor management path available on supported hardware platforms. It enables an administrator to reach the appliance independently of the BIG-IP host management interface and obtain the access needed to correct host-level network settings. A direct serial-console session is likewise independent of IP networking and provides local command-line access to BIG-IP. From either access path, the specialist can log in and modify the management address, resolving the duplicate-address condition without relying on the unavailable management GUI. F5 documents serial-console access as a recovery and local-administration mechanism, and documents AOM as an independent hardware-management capability on applicable appliances. See [F5 BIG-IP System Installation Essentials](#) and [F5 AOM information](#).

QUESTION NO: 51

-- Exhibit --

Custom ☐

General Configuration

Profile Name	avr_example		
Partition / Path	Common		
Parent Profile	analytics v		
Profile Description			
Statistics Logging Type	☒ Internal ☐ External	☒	
Traffic Capturing Logging Type	☐ Internal ☐ External	☒	
SMTP Configuration	None (Note: Setting can be changed only through the default analytics profile.)		
Notification Type	☒ Syslog ☐ SNMP ☐ E-mail	☒ ☐ ☐	
Trust XFF	☒ Enable ☐		
Transaction Sampling Ratio	Sample all transactions (Note: Setting can be changed only through the default analytics profile.)		

Included Objects

☐	Name	Destination	Service Port	Partition / Path
Virtual Servers				
No records to display.				
Add... Delete				

Custom ☐

Statistics Gathering Configuration

Collected Metrics	☒ Server Latency ☐ Page Load Time ☒ Throughput ☐ User Sessions	☒
Collected Entities	☐ URLs ☐ Countries ☐ Client IP Addresses ☒ Response Codes ☐ User Agents ☒ Methods	☐ ☐ ☐ ☒ ☐ ☒

Alerts and Notifications Configuration

Add New Rule
Alert when is Trans/sec, for seconds in

☐	Rule	Edit
☐	Alert when Average Server Latency is below 50 ms for 300 seconds in an Application.	Edit

Note: Changes you make might take up to 10 minutes to be reflected in the charts.

-- Exhibit --

Refer to the exhibit.

An LTM Specialist sets up AVR alerts and notifications for a specific virtual server if the server latency exceeds 50ms. The LTM Specialist simulates a fault so that the server latency is consistently exceeding the 50ms threshold; however, no alerts are being received.

Which configuration should the LTM Specialist modify to achieve the expected results?

- A. The rule should be adjusted to trigger when server latency is above 50ms.
- B. SNMP alerting should be enabled to allow e-mail to be sent to the support team.
- C. User Agents needs to be enabled to ensure the correct information is collected to trigger the alert.
- D. The metric "Page Load Time" needs to be enabled to ensure that the correct information is collected.

ANSWER: A

Explanation:

The rule should be adjusted to trigger when server latency is above 50ms. An AVR alert rule evaluates the selected metric using its configured comparison operator and threshold. Because the operational requirement is to notify the team when latency *exceeds* 50 ms, the condition must represent a value greater than 50 ms. With that comparison in place, sustained server-latency samples above the threshold meet the alert condition and AVR can generate the configured notification event for the selected virtual server.

Server latency is an application-performance measurement collected by AVR and represents time associated with the server-side portion of processing. It is therefore appropriate for detecting degradation introduced by a backend-server fault. The threshold is not itself sufficient: the rule's direction must match the desired failure condition. Configuring the condition as "above 50 ms" ensures that an increase in latency, rather than a decrease or an unrelated metric state, activates the alert. F5 documents AVR alerting and the available alert-rule metric conditions in the [BIG-IP Analytics alert configuration documentation](#). Additional context on Analytics metrics, including server-side latency reporting, is available in the [BIG-IP Analytics overview](#).

QUESTION NO: 52

-- Exhibit --

LTM device statistics

Search Reset Search

Status	Virtual Server	Partition / Path	Details	Bits		Packets		Connections		
				In	Out	In	Out	Current	Maximum	Total
☒	VS_HTTP	Common	View	283.9K	2.4M	391	544	0	5	55

Search Reset Search

Status	Pool/Member	Partition / Path	Bits		Packets		Connections		
			In	Out	In	Out	Current	Maximum	Total
☒	Pool_HTTP	Common	193.9K	2.4M	284	347	0	5	55
☒	-- 172.16.20.1:80	Common	103.4K	1.5M	163	205	0	1	19
☒	-- 172.16.20.2:80	Common	90.1K	872.4K	120	141	0	2	18
☒	-- 172.16.20.3:80	Common	416	0	1	0	0	2	19

-- Exhibit --

Refer to the exhibit.

An LTM Specialist is investigating intermittent page load issues being reported by users.

What should the LTM Specialist do to resolve the issue?

- A. Remove HTTP monitor on the pool.
- B. Assign an HTTP monitor to the pool.
- C. Select least connections load balancing method on virtual server.
- D. Remove least connections load balancing method on virtual server.

ANSWER: B

Explanation:

Assign an HTTP monitor to the pool. An HTTP health monitor actively validates that each pool member can provide the required web service, rather than merely assuming that a server is usable because its network address is reachable. The monitor periodically connects to the configured HTTP service and evaluates the response against its configured receive string or expected response behavior. If a member stops serving the application correctly, BIG-IP marks it unavailable and removes it from eligible load-balancing selections. When the application recovers and again satisfies the monitor, BIG-IP can return that member to service automatically.

This is especially important for intermittent page-load failures because a server can remain reachable at the TCP/IP level while its web service, application process, or HTTP response path is unhealthy. Associating the HTTP monitor with the pool gives BIG-IP application-aware health information and prevents requests from being sent to members that cannot successfully serve the monitored content. Configure the monitor with an HTTP request and expected response that

represent a meaningful application health check, ideally using a lightweight URL that depends on the required web service. F5 documents monitor behavior and pool-member availability in the [BIG-IP Local Traffic Management Monitors Reference](#) and its [health monitor troubleshooting guidance](#).

QUESTION NO: 53

What do the following iRule commands do when they are used in the same iRule?

```
set hsl [HSL::open -proto UDP -pool syslog_server_pool]
```

```
HSL::send $hsl " < 190 > [HTTP::host] from [whereis [IP::client_addr] country continent state city zip] , IP: [IP::client_addr] "
```

- A. The commands set up a high-speed logging connection and then send the geographical database to the server.
- B. The commands set up a high-speed logging connection and then send the host header and client geographical detail to the connection.
- C. The commands set up a high-speed logging connection and then send the host header, HTTP payload, and client geographical detail to the connection.
- D. The commands set up a high-speed logging connection to the LTM device and then send the host header and client geographical detail to the connection.

ANSWER: B

Explanation:

The commands set up a high-speed logging connection and then send the host header and client geographical detail to the connection. `HSL::open` creates an HSL publisher handle and stores it in the `hsl` variable. The command explicitly selects UDP as the transport and uses the pool named `syslog_server_pool`, so the pool members act as the destinations for the log stream. That handle is then supplied to `HSL::send`, which sends the constructed message through the already-open HSL connection.

The message includes `HTTP::host`, which inserts the HTTP Host header value. It also obtains the source address through `IP::client_addr` and passes that address to `whereis`. The requested `whereis` fields—country, continent, state, city, and ZIP code—supply the client geographic information included in the log message. Finally, the client IP address is appended directly. The leading priority value is formatted for a syslog-style message, while the HSL connection provides the high-speed UDP delivery path to the configured logging pool. F5 documents the handle-based HSL workflow in [HSL iRules reference](#) and the geographic lookup fields in the [whereis iRules reference](#).

QUESTION NO: 54

A BIG-IP Administrator needs to view the CPU utilization of a particular Virtual Server. Which section of the Configuration Utility should the administrator use for this purpose?

- A. Statistics > Module Statistics > Local Traffic > Virtual Addresses
- B. Statistics > Module Statistics > Traffic Summary
- C. Statistics > Analytics > Process CPU Utilization
- D. Statistics > Module Statistics > Local Traffic > Virtual Servers

ANSWER: D

Explanation:

Statistics > Module Statistics > Local Traffic > Virtual Servers is the appropriate Configuration utility location because it presents statistics at the individual virtual-server object level. An administrator can select the required virtual server and view its operational traffic and resource-related counters, including the CPU information associated with processing traffic for that

virtual server. This per-object view is important when the goal is to identify the behavior or load of one specific listener rather than assess system-wide utilization. The virtual-server statistics page also provides a focused operational context for correlating CPU activity with connection, throughput, packet, and availability statistics for the same virtual server. F5 documents virtual-server monitoring and statistics as part of Local Traffic Manager monitoring, where administrators use the Virtual Servers statistics view to inspect individual virtual-server performance. See F5's [Monitoring virtual servers documentation](#) and the [BIG-IP LTM virtual server reference](#) for the virtual-server object and its statistics context.