

DUMPS ARENA

Application Delivery Fundamentals

F5 771-101

Version Demo

Total Demo Questions: 17

Total Premium Questions: 178

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Application Delivery Fundamentals	89
Topic 2, Application Delivery Technologies	46
Topic 3, Application Delivery Security	37
Topic 4, Application Delivery Optimization	6
Total	178

QUESTION NO: 1

A BIG-IP high-availability pair uses floating traffic objects that must move to the peer device after a traffic-group failover.

Which two objects can be associated with a traffic group? (Choose two)

- A. A virtual address
- B. A floating self IP address
- C. A pool member
- D. A VLAN
- E. A management IP address

ANSWER: A B

Explanation:

A virtual address and **a floating self IP address** are correct. A traffic group defines the active device for floating traffic objects. When the traffic group fails over, the newly active device takes ownership of the associated virtual addresses and floating self IP addresses.

A pool member is an application destination and is not floated between devices through a traffic group. A VLAN is a local network object, while a management IP address is used for device administration rather than traffic-group ownership.

QUESTION NO: 2

A pool contains three application servers. One server has twice the processing capacity of each of the other servers, and the administrator wants it to receive approximately twice as many new connections during normal operation.

Which load-balancing method should be configured?

- A. Round Robin
- B. Least Connections (member)
- C. Ratio (member)
- D. Ratio Least Connections (member)

ANSWER: C

Explanation:

Ratio (member) is correct because it distributes new connections according to a configured ratio for each pool member. The higher-capacity server can be assigned a ratio value twice that of the other members, causing it to receive a proportionally larger share of traffic.

Round Robin distributes connections evenly without considering the configured capacity difference. Least Connections uses the current number of active connections rather than a fixed capacity ratio. Ratio Least Connections combines ratio weighting with active connection counts, which is unnecessary when the requirement is a static proportional distribution.

QUESTION NO: 3

What is one way that an application delivery controller lowers the CPU-load on web servers?

- A. use SSL termination at the application delivery controller
- B. make the application delivery controller function as a Firewall

- C. make the application delivery controller a member of both external and internal VLANs
- D. use cookie persistence on the application delivery controller

ANSWER: A

Explanation:

Using SSL termination at the application delivery controller lowers web-server CPU load by offloading the computationally intensive TLS/SSL work from the servers. Establishing secure client connections requires cryptographic operations such as certificate processing, key exchange, encryption, decryption, and message-integrity verification. When the application delivery controller terminates the client-side TLS connection, it performs these operations centrally and forwards the request to the selected web server, commonly over an internal connection that is either clear text or re-encrypted according to the organization's security requirements. This lets the web servers devote a greater share of their CPU capacity to application execution and serving content instead of repeatedly performing transport-security processing for every client connection. F5 BIG-IP's Client SSL profile is specifically designed to enable SSL/TLS termination on the virtual server, while the Server SSL profile can be used when encryption is also required from BIG-IP to the pool member. This SSL offload architecture is a core application-delivery function and can improve server scalability, especially for applications with high volumes of HTTPS traffic. See the [F5 BIG-IP SSL documentation](#) and the [F5 TLS/SSL deployment guidance](#).

QUESTION NO: 4

What is the correct procedure to comply with the recommendation?

- A. Download the product version image from ihealth f5 com
- B. Download the product version image from support 6 com
- C. Download the product version image from dovcentral f5.com
- D. Download the product version image from downloads f5.com

ANSWER: D

Explanation:

Download the product version image from downloads f5.com is the correct procedure because F5 distributes authorized BIG-IP software releases, hotfixes, engineering hotfixes, and associated image files through its official software-download service at downloads.f5.com. This portal is tied to F5 account authentication and product entitlements, helping ensure that the selected image is a legitimate F5 release appropriate for the licensed product and platform. Administrators can use the download service to select the required BIG-IP product version and obtain the installation image needed for an upgrade, installation, or recovery workflow. Obtaining software from the official F5 download location also supports operational integrity: the downloaded image can be checked against the checksum or signature information F5 publishes with the release before it is installed. F5's product documentation directs customers to use the F5 Downloads site when obtaining BIG-IP software and related files; see the [F5 Technical Documentation](#) portal for version-specific installation and upgrade guidance. Using the authorized download service ensures the organization follows F5's supported software-acquisition process.

QUESTION NO: 5

A BIG-IP Administrator is reviewing a TCP packet capture for a new client connection.

Which two events occur during the normal TCP three-way handshake? (Choose two)

- A. The client sends a SYN segment
- B. The server responds with a SYN-ACK segment
- C. The server sends a FIN segment
- D. The client receives an ICMP Echo Reply

E. The client sends HTTP application data before receiving a SYN-ACK

ANSWER: A B

Explanation:

The client sends a SYN segment and the server responds with a SYN-ACK segment are correct. The client initiates a TCP connection with SYN, and the server acknowledges that request while sending its own SYN. The client then completes the handshake with an ACK.

Application data is normally sent after the connection is established, not as a required part of the three-way handshake. A FIN segment is used to begin orderly connection termination, and an ICMP Echo Reply is unrelated to TCP session establishment.

QUESTION NO: 6

What are two examples of failover capabilities of BIG-IP? (Choose two)

- A. failing over based on an over temperature alarm
- B. failing over serial cable based on electric failure code
- C. II failing over network connection based on heartbeat detection
- D. fading over network connection based on SNMP error
- E. failing over serial cable based on voltage detection

ANSWER: C E

Explanation:

BIG-IP high availability supports peer-health detection through both network-based and serial-cable-based failover paths. **II failing over network connection based on heartbeat detection** describes network failover: devices in a device service cluster exchange heartbeat and failover information across the configured failover network, allowing the peer to determine that its partner is unavailable and assume the active role when appropriate. This mechanism provides a redundant, IP-based communication path for high-availability operation.

failing over serial cable based on voltage detection represents serial failover. A direct serial connection supplies an additional, independent communication path between the BIG-IP devices. The serial interface's electrical signaling and loss of serial connectivity can be used in peer-failure detection, which is particularly valuable when the normal network failover path is unavailable. Using both network and serial paths improves the reliability of failover decisions because device communication is not dependent on only one transport mechanism.

F5 documents network and serial failover as the primary failover communication methods for a BIG-IP device service cluster. See the [BIG-IP Device Service Clustering Administration guide](#) and F5's [Device Service Clustering documentation collection](#).

QUESTION NO: 7

A virtual server must accept HTTPS connections from clients, decrypt the traffic for HTTP processing on BIG-IP, and then establish a new encrypted HTTPS connection to the pool members.

Which two profiles are required? (Choose two)

- A. A Client SSL profile
- B. A Server SSL profile
- C. An HTTP compression profile
- D. A persistence profile

ANSWER: A B

Explanation:

A Client SSL profile is required to terminate and decrypt the client-side TLS connection. It contains the certificate, private key, and client-side TLS settings used by BIG-IP.

A Server SSL profile is required when BIG-IP must create a separate TLS connection to the pool member. An HTTP profile can enable HTTP processing after decryption, but it does not perform either client-side or server-side TLS processing.

QUESTION NO: 8

BIG-IP Administrator runs the command `ITconfig` from the BIG-IP command line and sees a high number of collisions. What is the likely cause?

- A. There is a duplicate IP address on the network
- B. A crossover cable has been used
- C. The interface is set to half duplex
- D. Spanning Tree Protocol is blocking the connection

ANSWER: C

Explanation:

The interface is set to half duplex is the likely cause. Ethernet collisions are a normal consequence of the CSMA/CD access method used by half-duplex Ethernet: because the transmit and receive paths share a collision domain, two devices can begin transmitting at nearly the same time. The interface detects this condition, stops transmission, and retries after a backoff interval. Therefore, a sustained or high collision counter on a BIG-IP interface is a strong indicator that the link is operating in half-duplex mode, frequently as the result of an incorrect speed/duplex setting or unsuccessful autonegotiation with the connected switch port.

On a properly operating full-duplex Ethernet link, sending and receiving occur simultaneously on separate paths and collisions should not be present. An administrator should inspect the BIG-IP interface's negotiated link state and compare its speed and duplex configuration with the directly connected switch port. BIG-IP exposes interface configuration and operational details through the `net interface` TMOS object; see the F5 [tmsh net interface reference](#). Cisco's Ethernet troubleshooting guidance also describes collisions as a half-duplex Ethernet behavior and recommends verifying duplex agreement between link peers: [Troubleshooting Ethernet](#).

QUESTION NO: 9

A remote-access policy must verify a user identity and then determine whether that user is entitled to access a protected application.

Which two AAA functions are required? (Choose two)

- A. Authentication
- B. Authorization
- C. Accounting
- D. Availability
- E. Acceleration

ANSWER: A B

Explanation:

Authentication verifies the identity of a user, such as validating a username and password against an LDAP or Active Directory server. Authorization determines what an authenticated user is allowed to access, such as an application, network resource, or specific role-based entitlement.

Accounting records session or resource-usage information after access activity occurs. Availability is an application-delivery objective, but it is not an AAA function.

QUESTION NO: 10

A BIG-IP high-availability pair is configured to provide application availability during a device failure. Which two statements correctly describe configuration synchronization and connection mirroring? (Choose two)

- A. Configuration synchronization shares selected configuration objects with the peer
- B. Connection mirroring shares eligible connection state with the peer
- C. Configuration synchronization sends all application traffic to both devices
- D. Connection mirroring automatically balances new connections across both devices
- E. Configuration synchronization replaces the need for a device group

ANSWER: A B

Explanation:

Configuration synchronization distributes selected BIG-IP configuration changes between members of a device group. It helps ensure that the peer has the virtual servers, pools, profiles, and other synchronized objects needed after failover.

Connection mirroring shares connection-state information with the peer so that eligible active connections can continue with less disruption after failover. Configuration synchronization alone does not mirror active connection state, and connection mirroring does not replace configuration synchronization.

QUESTION NO: 11

In which scenario does the client act as a server?

- A. During an active FTP session
- B. When an SMTP connection
- C. During an SMTP connection
- D. When browsing websites

ANSWER: A

Explanation:

During an active FTP session, the client temporarily takes the server role for the FTP data connection. FTP uses separate control and data connections. After the client establishes the control connection to the FTP server, it sends a PORT command in active mode identifying an address and port on which it will listen for the data channel. The FTP server then initiates the data connection back to that listening client port. Because the client is accepting an inbound connection for that data transfer, it is functioning as a server for that connection, even though it remains the FTP client from the overall application perspective.

This behavior is important in application delivery and firewall design because the inbound, server-initiated data channel may be blocked by client-side firewalls or NAT devices unless they understand and dynamically permit the FTP active-mode flow. F5 FTP profiles can inspect FTP control traffic and assist with the related data connection handling. The FTP specification defines the PORT command and the server-initiated data connection behavior in [RFC 959](#). F5 documentation also describes the FTP profile and its handling of FTP connections in the [F5 CloudDocs FTP profile documentation](#).

QUESTION NO: 12

A user receives a browser certificate warning while accessing an HTTPS virtual server. The administrator confirms that BIG-IP presents a certificate signed by a public certificate authority.

Which two certificate conditions must the browser validate before it can trust the certificate for the requested site? (Choose two)

- A. The certificate must be within its validity period
- B. The requested hostname must match the certificate identity
- C. The virtual server must use TCP port 443
- D. The browser must receive the server private key
- E. The certificate serial number must match a trusted root certificate

ANSWER: A B

Explanation:

The certificate must be within its validity period and the requested hostname must match the certificate identity are correct. A browser checks that the certificate is not expired or not yet valid and verifies that the requested DNS name matches a Subject Alternative Name or other valid certificate name.

Using TCP port 443 does not establish certificate trust. The server private key must remain private and is not sent to the browser. A certificate does not need to have the same serial number as a browser root certificate; it must instead chain to a trusted certificate authority.

QUESTION NO: 13

A BIG-IP Administrator has a cluster of devices.

What should the administrator do after creating a new Virtual Server on device 1?

- A. synchronize the settings of device 1 to the group
- B. create a new cluster on device 1
- C. create the new virtual server on device 2
- D. synchronize the settings of the group to device 1

ANSWER: A

Explanation:

After creating a new virtual server on device 1, the administrator should **synchronize the settings of device 1 to the group**. In a BIG-IP device service clustering (DSC) deployment, virtual servers are configuration objects that must be shared with peer devices in the same sync-failover device group. Synchronizing from the device where the change was made publishes that device's updated configuration to the other group members. This ensures that every device has the same virtual-server definition, associated pool references, profiles, policies, and other relevant configuration data. Maintaining synchronized configurations is essential for high availability: if traffic must fail over to another device, that device needs the virtual server already present and configured identically to continue serving applications. BIG-IP can be configured for

manual or automatic synchronization, but the required action after a local configuration change is to push the updated configuration from the changed device to its device group. F5 documents configuration synchronization as the mechanism by which devices in a sync-failover group maintain a common configuration. See the [F5 BIG-IP Configuration Synchronization documentation](#) and the [F5 DSC and high-availability training material](#).

QUESTION NO: 14

An administrator needs to lower the amount of new TCP connections to help with network congest?

How should the administrator accomplish this?

- A. Use HTTP keep-alive
- B. set a larger receive window
- C. Enable delayed ACKs
- D. use an indefinite timeout

ANSWER: A

Explanation:

Use HTTP keep-alive is correct because it allows multiple HTTP requests and responses to use an already-established TCP connection rather than requiring a separate TCP three-way handshake and connection teardown for each object or request. Reusing persistent connections directly reduces the number of new TCP connections clients must create, which lowers connection-setup traffic and associated processing on clients, servers, and the BIG-IP system. This is especially beneficial for web applications whose pages require many resources, such as images, style sheets, scripts, and API calls. On BIG-IP, the HTTP profile controls HTTP connection persistence behavior, including the Keep-Alive setting. When keep-alive is enabled and the client and server support persistent HTTP connections, BIG-IP can retain the connection for subsequent requests until the configured idle timeout or connection limits are reached. This improves efficiency by amortizing the TCP setup cost across multiple application transactions and can help reduce network congestion caused by frequent connection establishment. F5 documents the HTTP profile's keep-alive controls in its [HTTP::keepalive reference](#) and describes HTTP profile configuration in the [BIG-IP Local Traffic Manager Profiles Reference](#).

QUESTION NO: 15

Which security model provides protection from unknown attacks?

- A. ACL Security model
- B. closed security model
- C. Negative security model
- D. Positive security model

ANSWER: D

Explanation:

The Positive security model provides protection from unknown attacks because it defines an explicit allowlist of expected, valid application behavior. Rather than needing a previously identified attack signature, the security policy permits only known-good requests, such as approved URLs, methods, parameters, parameter data types, input lengths, character sets, and file types. Any request that falls outside that established application profile can be blocked or flagged, even when it represents a newly developed exploit technique or a zero-day attack.

This approach is particularly effective for web application firewalls because many attacks rely on sending unexpected input, invoking unapproved endpoints, or using parameter values that do not conform to the application's intended design. By enforcing what the application is supposed to accept, a positive model reduces reliance on continually updating a list of malicious patterns. OWASP describes allowlisting as the preferred input-validation approach when practical, because it

accepts only data matching defined valid criteria. F5 WAF policy concepts likewise support application-specific enforcement through configured URL, parameter, and other entity definitions. See the [OWASP Input Validation Cheat Sheet](#) and [F5 WAF Declarative Policy documentation](#).

QUESTION NO: 16

Which two capabilities are provided by a full proxy architecture? (Choose two)

- A. Independently manage client-side and server-side TCP connections
- B. Modify HTTP headers before forwarding a request
- C. Forward return traffic without receiving it from the server
- D. Eliminate the need for a virtual server
- E. Operate without maintaining connection state


ANSWER: A B

Explanation:

A full proxy terminates the client-side connection and establishes a separate connection to the pool member. This allows BIG-IP to apply independent TCP behavior on each side of the connection and to perform application-layer processing, such as modifying HTTP headers.

A packet-forwarding architecture forwards packets without creating separate client-side and server-side connections. It can provide lower latency, but it does not provide the same content-aware processing capabilities.

QUESTION NO: 17 - (SIMULATION)



An administrator receives an error message when attempting to create a self-IP address on a BIG-IP appliance. Which address should the administrator use to resolve the issue?

- A. 2001 0db8.85a3 0000 0000 8a2e 0370 7334
- B. 2001:db8:85a3 Bd3.Ba2e.370 7334
- C. 2001:0db8:85i3:0000:0000: : 8a2i:0370:7334

ANSWER: See the explanation for the answer

Explanation:

Answer: A — use:

This is a valid IPv6 address: it has eight colon-separated hexadecimal groups, and each group contains only valid hexadecimal characters (0–9 and a–f).

The attempted address shown in the error includes `h` characters (for example, `85h3` and `8a2h`), which are not valid in IPv6 hexadecimal notation. The other choices also contain invalid characters or invalid separators.