PassQueen
a.k.a dumpsqueen.com

Application Delivery Fundamentals

F5 101

Version Demo

Total Demo Questions: 71

Total Premium Questions: 711

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Networking Fundamentals	60
Topic 2, Protocol Fundamentals	59
Topic 3, Application Delivery Fundamentals	38
Topic 4, Security Fundamentals	111
Topic 5, F5 Technology Fundamentals	441
Topic 6, Mix Questions	2
Total	711

QUESTION NO: 1

Which three of these are the potential ending options for branches in the Visual Policy Editor? (Choose three.)

A. Reject

B. Fallback

C. Allow

D. Deny

E. Limit

F. Redirect

ANSWER: C D F

Explanation:

In an F5 BIG-IP Access Policy Manager access policy, the Visual Policy Editor uses ending actions to determine the final result when a user reaches the end of a policy branch. **Allow** ends the branch by granting the session access to the protected resource. **Deny** ends the branch by refusing access, which is used when policy conditions do not authorize the user or session. **Redirect** ends the branch by sending the user to a specified URL, allowing the policy to direct the user to an alternate destination or workflow rather than granting or denying access directly.

These ending actions are placed at the conclusion of a branch after the relevant authentication, endpoint, identity, or authorization checks have been evaluated. They provide the definitive access-policy outcome for the matching path, so a policy can produce different outcomes based on the branch selected by preceding agents and branch rules. F5 documents the Visual Policy Editor as the interface for constructing access-policy flows and configuring branch outcomes, including terminal access decisions. See the F5 BIG-IP Access Policy Manager Visual Policy Editor documentation and the F5 guidance for ending an access policy.

QUESTION NO: 2

Which two functions of AAA does Security Assertion Markup Language (SAML) identity Provider (IdP) offer (Choose two)?

A. measure usage

B. provide authentication

C. provide availability

D. provide session auditing

E. filter access

ANSWER: B E

Explanation:

SAML identity providers support the authentication and authorization portions of the AAA model. **provide authentication** is correct because the IdP establishes the user's identity by authenticating the user through its configured authentication method. Following successful authentication, the IdP creates and signs a SAML assertion for the service provider (SP). This assertion communicates that the user has authenticated and supplies identity information in a trusted, standards-based format.

filter access represents authorization. A SAML assertion can carry attributes, roles, group membership, or authorization-decision information that an SP or access-policy implementation uses to determine which protected resources the authenticated user may access. In an F5 BIG-IP Access Policy Manager deployment, the IdP and access policy can therefore provide the identity and entitlement context needed to apply access controls to an application. SAML's assertion framework defines authentication statements and authorization decision statements, allowing identity and access decisions

to be conveyed between federation participants. For further detail, see the OASIS SAML 2.0 Core specification and F5's BIG-IP APM Authentication and Single Sign-On documentation.

QUESTION NO: 3 - (SIMULATION)

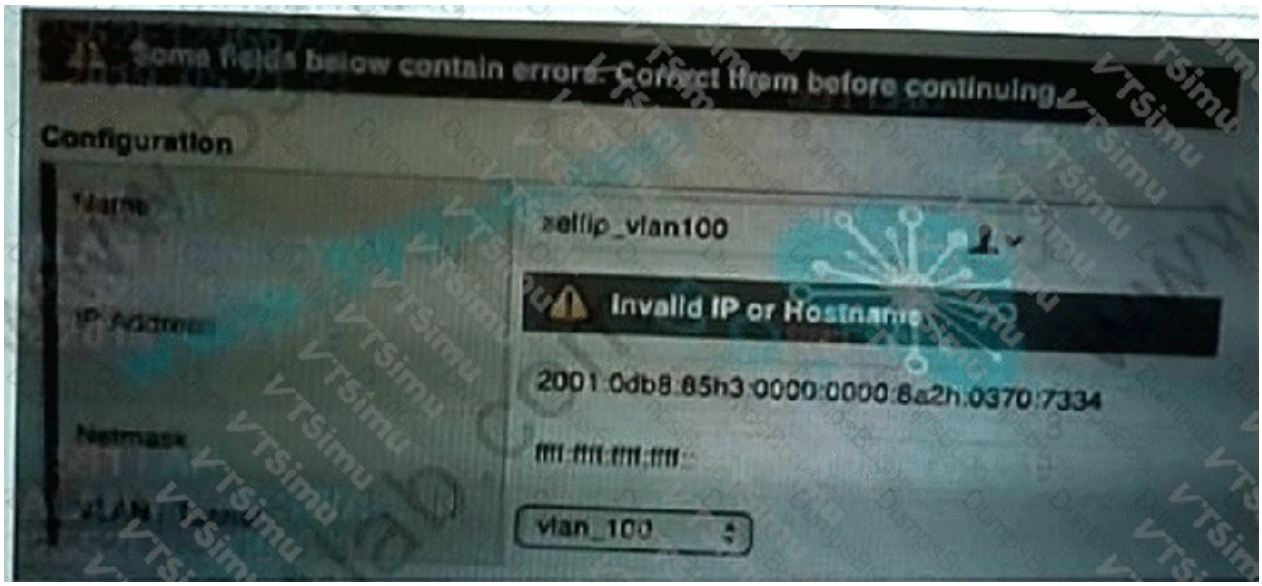

An administrate receives an error message when attempting to create a self-IP address on a BIG-IP appliance Which address should the administrator use to resolve the issue?

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: B — `2001:db8:85a3::8a2e:0370:7334`

This is a valid IPv6 address. The `::` abbreviation represents consecutive all-zero 16-bit groups, so the expanded form is:

An IPv6 hextet may contain only hexadecimal characters (`0-9` and `a-f`). The attempted address in the image contains invalid characters such as `h`; choice C similarly contains invalid `i` characters. Use option B when creating the self IP.

QUESTION NO: 4

TI SISSL prides which functionary when a web browser connects to web site using the HTTPS protocol?

A. delivers secure communication over a network connection

B. prevents the web server from being compromised

C. protects network packets from interception by intermedianes

D. stop the web browser from automatically downloading malware

ANSWER: A

Explanation:

delivers secure communication over a network connection is correct because HTTPS uses TLS (formerly commonly called SSL) to secure the connection between a browser and a web server. During the TLS handshake, the parties negotiate security parameters, authenticate the server through its certificate, and establish cryptographic keys for the session. Those keys protect application data sent over the HTTPS connection, providing confidentiality so the content is not readable by unauthorized parties, integrity so modifications can be detected, and authentication of the server identity when certificate validation succeeds.

HTTPS is therefore not simply a different web protocol; it is HTTP carried within a TLS-protected channel. The browser's request, response content, credentials, cookies, and other HTTP data are protected while in transit across networks. TLS is specifically defined as a protocol for securing communications between client and server endpoints, and HTTPS is defined as the use of HTTP over TLS. See the IETF's TLS 1.3 specification (RFC 8446) and HTTP Semantics specification (RFC 9110) for the protocol foundations.

QUESTION NO: 5

What is the primary purpose of a VLAN on a BIG-IP system?

A. To define a Layer 2 network segment for traffic interfaces

B. To select the least-loaded pool member

C. To store SSL certificates

D. To maintain client persistence records

E. To perform DNS resolution

ANSWER: A

Explanation:

To define a Layer 2 network segment for traffic interfaces is correct. On BIG-IP, VLANs associate one or more interfaces with a Layer 2 broadcast domain. Self IP addresses are assigned to VLANs so BIG-IP can communicate on the relevant network segments.

A VLAN is not a load-balancing method, a persistence mechanism, or an SSL certificate container. F5 documents VLAN configuration in the BIG-IP VLAN reference.

QUESTION NO: 6

The ARX saves customers time, money and frustration through a stub-based system that makes a slight modification to each file in order to more efficiently sort and store end user data.

A. True

B. False

ANSWER: B

Explanation:

False is correct because F5 ARX is a file-virtualization platform, not a stub-based hierarchical-storage system. ARX presents a global, location-independent namespace to SMB/CIFS and NFS clients while it transparently directs file requests to the appropriate physical file server or storage tier. Its policy engine can analyze file attributes and move or place files among back-end storage resources, enabling organizations to use capacity and performance tiers more efficiently without requiring users to change the paths they use to access data.

Crucially, ARX maintains the virtual namespace and routing information rather than modifying every migrated user file to insert a stub or placeholder. This architecture permits file movement and storage consolidation while preserving transparent client access through the ARX device. The benefit described in the question—reducing administrative effort and improving storage efficiency—is consistent with ARX, but attributing that function to per-file stub creation is not. F5 describes ARX as

providing file virtualization, policy-based data management, and transparent data movement across heterogeneous storage systems. See the F5 ARX Deployment Guide and F5's ARX product information.

QUESTION NO: 7

Which four of these scenarios will benefit from F5's WAN Optimization Module? (Choose four.)

A. An international organization with data centers in different countries.

B. An organization that does not want to rely on using tape backup.

C. An organization with one site but hundreds of Web servers.

D. An organization whose users create extremely large files.

E. An organization that expects their Web site usage to double in the next year.

F. An organization attempting to lower costs by reducing the number of data centers.

ANSWER: A B D F

Explanation:

F5's WAN Optimization Module is intended to improve application and data-transfer performance across constrained or high-latency WAN links. An international organization with data centers in different countries benefits because optimization reduces the WAN effects experienced by users and applications communicating between geographically separated locations. An organization that does not want to rely on using tape backup benefits because WAN optimization can make network-based backup, replication, and data movement more practical by reducing transferred data and accelerating WAN traffic.

An organization whose users create extremely large files also benefits from WAN data reduction, compression, and transport optimization, which can substantially improve the efficiency of repeated or large transfers between sites. Finally, an organization attempting to lower costs by reducing the number of data centers can centralize applications and data while using WAN optimization to help remote offices retain acceptable access performance. These are core use cases for WAN optimization: accelerating remote access, enabling data replication and backup over WAN links, and supporting data-center consolidation initiatives. F5 describes WAN optimization as improving application performance and reducing bandwidth consumption for traffic between distributed locations. See the F5 WAN Optimization Module deployment guide and F5 BIG-IP WOM documentation.

QUESTION NO: 8

The ARX can see _____ when a data modification takes place and will cue that file to be migrated back to the primary tier.

A. In real time

B. Nightly

C. Weekly

D. At the time of a system scan

E. When scheduled by administrator

ANSWER: A

Explanation:

In real time is correct. F5 ARX operates in the client access path for file-serving environments, presenting a global namespace while directing file requests to the appropriate back-end storage. Because ARX processes the relevant file-access traffic, it can detect a modification as it occurs rather than relying on a later inventory operation or a periodic

administrative task. This immediate visibility supports automated tiering behavior: when a file that resides on a lower-cost or secondary tier is modified, ARX can cue the file for migration back to the primary tier. Returning an actively changed file to the primary tier helps ensure that current, frequently accessed content is kept on the storage tier intended for active workloads, while less active content can remain on capacity-oriented storage. This behavior is part of policy-driven file virtualization and storage-tier management, not a batch-only process. F5 describes ARX as a file virtualization platform that provides a global namespace and policy-based control of file placement across storage resources. See the F5 ARX data sheet and the F5 ARX knowledge center for product documentation and support resources.

QUESTION NO: 9

An administrator is planning on solving latency issues by placing the backup data center in a neighboring city to the primary data center. Why isn't this effective solution?

A. It does not provide an acceptable disaster recovery solution.

B. It won't improve latency between the data centers.

C. It won't improve the backup time from the primary data center to the backup data center.

D. Users near the data centers will find their connections going back and forth between the two site

E. It costs more to have the data centers close together than further apart.

ANSWER: A

Explanation:

It does not provide an acceptable disaster recovery solution. is correct because a backup site must be sufficiently isolated from the primary site to avoid a single regional event affecting both locations. Although placing a secondary data center in a neighboring city can reduce replication latency and support a lower recovery-point objective, close geographic proximity can expose both sites to the same disaster domain. For example, a major weather event, earthquake, flood, prolonged regional power failure, telecommunications outage, or other area-wide disruption could make both the primary and backup locations unavailable.

Effective disaster-recovery planning balances replication performance against the need for geographic and infrastructure diversity. The recovery site should be selected through a risk assessment that considers correlated threats, independent power and network providers, and the organization’s recovery objectives. A nearby site may be appropriate for high-availability designs, but it does not by itself provide the separation expected of a resilient disaster-recovery strategy. NIST contingency-planning guidance emphasizes identifying alternate processing facilities that can support recovery when the primary facility is disrupted. See NIST SP 800-34 Rev. 1, Contingency Planning Guide and F5 White Papers for disaster-recovery and application-availability architecture resources.

QUESTION NO: 10

HTTPs traffic is not working properly.

What is the likely reason for the issue?

A. 0. 0 0.0 0 is an invalid address in netstat.

B. 0. 0. 0. 0.0 80 should be in an active stale

C. The server is not listening on TCP 443

D. The server is not listening on UDP 80

ANSWER: C

Explanation:

The server is not listening on TCP 443. HTTPS is HTTP protected by TLS, and its well-known transport endpoint is TCP port 443. For a client to establish an HTTPS connection, the destination system must have an application or virtual service bound to TCP 443 and in a listening state. A listening socket on port 443 accepts the initial TCP connection, after which the TLS handshake and encrypted HTTP exchange can proceed. The displayed listener information shows availability for HTTP on TCP port 80, but it does not show a TCP listener for port 443. Consequently, HTTPS requests cannot reach a service that can accept and process them. In an F5 deployment, this same principle applies to the configured virtual server: the virtual server must listen on the expected HTTPS destination port and have an appropriate client SSL profile when BIG-IP is responsible for TLS termination. TCP port 443 is registered as the standard service port for HTTPS by the Internet Assigned Numbers Authority; see the IANA service-name and port-number registry. For BIG-IP SSL/TLS processing concepts and profile configuration, consult the F5 BIG-IP documentation.

QUESTION NO: 11

Which two processes are involved when BIG-IP systems issue traps? (Choose two.)

A. bigd

B. alertd

C. smtpd

D. syslogng

ANSWER: B D

Explanation:

BIG-IP trap generation involves **alertd** and **syslogng** (the syslog-ng logging service). The syslog-ng service receives, processes, and routes system log messages generated by BIG-IP components. These messages provide the event information from which alerts can be generated. The alertd daemon monitors the applicable alert conditions and is responsible for delivering configured alert notifications, including SNMP traps, to the defined SNMP trap destinations. Together, these services provide the event-to-notification path: system activity is recorded and handled by the logging subsystem, then alert processing generates the outbound trap notification when the event matches an alert definition or threshold. This distinction is useful operationally because logging and alert delivery are separate functions, even though they cooperate in the trap workflow. F5 documents SNMP configuration and trap destinations in its BIG-IP monitoring documentation, while its logging documentation describes syslog-ng as the system logging service. See the BIG-IP SNMP documentation and the BIG-IP logging documentation.

QUESTION NO: 12

In the following request, which portion represents a parameter name?

A. Yes

B. User

C. Week1

D. Financials

ANSWER: B

Explanation:

User is the parameter name. In an HTTP request URI, data passed in the query component is commonly structured as name/value pairs. The name appears to the left of an equals sign, while its associated value appears to the right. Thus, in a query-string element such as `User=Week1`, `User` identifies the input field or variable being supplied, and `Week1` is the value assigned to that field. Parameter names let the receiving web application identify what each supplied value represents—for example, a user identifier, search term, page number, or filter setting—before it processes the request. This distinction is important when examining HTTP traffic, configuring traffic-management policies, or troubleshooting application behavior on an F5 device, because matching on a parameter name targets the field itself rather than a particular submitted value. URI query syntax is defined broadly by RFC 3986, while the use of query parameters is application-specific. See RFC 3986, section 3.4: Query and MDN: Sending form data.

QUESTION NO: 13

Host A sends 10 TCP packets to Host

B.
All packets arrive at Host B quickly, but some arrive out of order.

What will Host B do?

A. ACK only packets that are in order

B. drop all packets and wait for arrival in order

C. drop any packets that arrive out of order

D. ACK all packets and place them in order

ANSWER: A

Explanation:

ACK only packets that are in order is the best answer under the fundamental TCP acknowledgment model. TCP is a reliable, ordered byte-stream protocol, so the receiving host uses TCP sequence numbers to identify where each received segment belongs in the stream. When a segment arrives ahead of a missing segment, the receiver does not advance its cumulative acknowledgment number beyond the missing data. Instead, it acknowledges the next byte it expects to receive in order. This lets the sender recognize that a gap exists and retransmit the missing data if necessary.

The receiver can retain out-of-order data in its receive buffer while waiting for the missing sequence range. Once that missing data arrives, the receiver can deliver the now-contiguous byte stream to the application and advance its acknowledgment accordingly. In implementations supporting Selective Acknowledgment, the receiver may also report noncontiguous blocks that it has received, but the fundamental cumulative ACK still identifies the next in-order byte expected. This behavior preserves TCP's ordered delivery guarantee. See RFC 793, Transmission Control Protocol and RFC 1122, TCP acknowledgment requirements.

QUESTION NO: 14

Which of the following user roles have access to make changes to security policies? (Choose two.)

A. Guest

B. Operator

C. Administrator

D. Web Application Security Editor

ANSWER: C D

Explanation:

Administrator and **Web Application Security Editor** have the permissions needed to change BIG-IP Application Security Manager security policies. The Administrator role has unrestricted administrative access to the system and its configuration, including Application Security Manager policy objects. This permits creating, editing, importing, exporting, applying, and otherwise managing security-policy configuration.

The Web Application Security Editor role is specifically intended for personnel who manage web application security. It grants access to Application Security Manager policy configuration, allowing the user to modify policy settings and the security controls associated with protected applications. This separation supports role-based administration: a security specialist can maintain web application protections without requiring full system-wide administrative authority.

F5 documents that BIG-IP uses role-based access control to assign permissions according to a user's operational responsibilities. Application Security Manager policy administration is therefore available to full administrators and to the dedicated web application security editing role. See the F5 BIG-IP Application Security Manager documentation and the BIG-IP user account administration guide for role and access-control guidance.

QUESTION NO: 15

Since F5 built GTM on the TMOS platform it can exist on the same BIGIP device as LTM:

A. True

B. False

ANSWER: A

Explanation:

True is correct. Global Traffic Manager (GTM), now generally referred to as the BIG-IP DNS module, is a software module that runs on the BIG-IP system's TMOS platform. Local Traffic Manager (LTM), now the BIG-IP Local Traffic Manager module, is also provisioned on TMOS. Consequently, a single BIG-IP device can have both modules licensed and provisioned, allowing it to perform local application traffic management as well as DNS and global traffic-management functions.

When deployed together, the modules share the BIG-IP platform while retaining their respective configuration objects and traffic-processing roles. The practical requirement is that the appliance or virtual edition has appropriate licenses and sufficient hardware resources for the modules that are provisioned. Administrators can then configure LTM virtual servers, pools, and profiles alongside BIG-IP DNS wide IPs, pools, and DNS services on the same system. F5 documentation identifies BIG-IP DNS as a TMOS software module and describes provisioning modules on a BIG-IP device, supporting this co-resident deployment model.

See the F5 BIG-IP DNS documentation and the F5 BIG-IP Local Traffic Manager documentation.

QUESTION NO: 16

An administrator needs to restore the BIG-IP configuration from a UCS file In which section in the BIG IP Configuration utility is this task performed?

A. System > Archives

B. System > Services > Backup

C. System > Backup

D. System > File Management

ANSWER: A

Explanation:

System > Archives is the correct location because the BIG-IP Configuration utility uses the Archives screen to manage BIG-IP archive files, including User Configuration Set (UCS) files. A UCS archive is a consolidated backup of the device configuration and related system files. From this area of the GUI, an administrator can upload a UCS archive to the BIG-IP system when necessary and then restore it, returning the configuration to the state captured in that archive. This workflow is especially useful for device replacement, recovery after configuration loss, or moving an existing configuration to a compatible BIG-IP system. Administrators should ensure that the UCS is appropriate for the target platform and software version, and should account for device-specific settings such as management addressing and licensing as part of the restoration plan. F5 documents UCS archives as the mechanism for backing up and restoring BIG-IP configuration data, while the Configuration utility's Archives section is the GUI location for archive operations. See F5's guidance on UCS archives and the BIG-IP documentation for managing configuration archives.

QUESTION NO: 17

Which three of these file types work well with HTTP compression? (Choose three.)

A. MP4 videos

B. Digital photos

C. Text files

D. Static HTML Web pages

E. CD quality songs

F. Microsoft Word documents

ANSWER: C D F

Explanation:

Text files, static HTML Web pages, and Microsoft Word documents work well with HTTP compression because they generally contain substantial repetitive textual and structural data. Compression algorithms used by HTTP, such as gzip, are especially effective when content includes recurring words, markup tags, whitespace, formatting structures, and other predictable byte patterns. Text files and static HTML pages are classic examples: they are usually delivered as readable, uncompressed text and can often be reduced significantly before transmission. Microsoft Word documents, particularly traditional document formats containing text and structured formatting data, can also benefit when their content is not already stored in a compressed container format.

On BIG-IP, an HTTP Compression profile applies compression to eligible HTTP responses based on settings such as the response content type, response size, and the client's `Accept-Encoding` request header. This reduces the number of bytes sent across the network, which can improve page and document delivery times while conserving bandwidth. The F5 HTTP Compression profile documentation describes the configurable content types and response conditions used for compression, while HTTP compression negotiation is defined by the HTTP content-coding standard. See F5 BIG-IP HTTP Compression profile reference and RFC 9110: Accept-Encoding.

QUESTION NO: 18

Which of the following statements are correct regarding Attack signatures? (Choose two.)

A. Attack signatures can apply to requests, responses, and parameters.

B. Attack signatures are the basis for positive security logic with the BIG-IP ASM System.

C. Any new Attack signature downloaded manually or automatically will be active and assigned directly to the security policy.

D. Individual Attack signatures cannot be assigned directly to the security policy. Only Attack signature sets can be applied to the security policy.

ANSWER: A D

Explanation:

Attack signatures can apply to requests, responses, and parameters because BIG-IP ASM examines HTTP traffic for known malicious patterns at several relevant inspection points. Request signatures identify attack content in client-supplied HTTP request data, including URLs, headers, and payloads. Parameter-level enforcement provides more granular inspection by applying relevant signatures to the values submitted for specific parameters. Response inspection can similarly identify sensitive-information leakage or response-based attack patterns when response checking is configured. This coverage lets a security policy detect known attack techniques across both inbound and outbound application traffic.

Attack signature sets are the policy-level mechanism used to apply signature protections. A set groups related signatures, commonly by technology, risk, or attack category, so the policy can consistently enable the appropriate collection of detections rather than requiring direct policy assignment of each individual signature. Administrators can tune signature behavior within the context of a policy and its enabled sets, which supports a controlled rollout of protections appropriate to the application. F5 documents attack-signature inspection and signature-set configuration in its BIG-IP ASM Attack Signatures documentation and provides signature-update guidance in F5 Knowledge Center article K12223.

QUESTION NO: 19

A site needs to terminate client HTTPS traffic at the BIG-IP and forward that traffic unencrypted. Which two are profile types that must be associated with such a virtual server? (Choose two.)

A. TCP

B. HTTP

C. HTTPS

D. ClientSSL

E. ServerSSL

ANSWER: A D

Explanation:

TCP and **ClientSSL** are the required profile types for this SSL-offload design. A Client SSL profile is attached to the client side of the virtual server and enables BIG-IP to perform the TLS handshake with the client, present the site certificate, and decrypt the incoming HTTPS payload. After that decryption occurs, BIG-IP has access to the clear-text application traffic and can send it to the pool member without encrypting the server-side connection.

A TCP profile provides the Layer 4 TCP handling used for the client connection and for the proxied server-side flow. HTTPS runs over TCP, so TCP connection processing is fundamental to accepting the client HTTPS session and forwarding the resulting HTTP data. BIG-IP supplies TCP profile functionality by default for an appropriate virtual-server configuration, but the TCP profile type is nevertheless part of the required traffic-processing path. No server-side TLS profile is used when the explicit requirement is to forward traffic unencrypted. For F5 configuration details, see the BIG-IP Client SSL profile reference and the BIG-IP TCP profile reference.

QUESTION NO: 20

Which HTTP response code ranges indicate an error condition? (Choose two.)

A. 1xx

B. 2xx

C. 3xx

D. 4xx

E. 5xx

ANSWER: D E

Explanation:

The HTTP status-code classes that indicate error conditions are **4xx** and **5xx**. A 4xx response identifies a client-error condition: the server received and understood the request but cannot fulfill it because of an issue associated with the request, such as missing authentication credentials, insufficient authorization, or a resource that does not exist. Common examples include 401 Unauthorized, 403 Forbidden, and 404 Not Found.

A 5xx response identifies a server-error condition. It means that the server encountered an unexpected condition or is otherwise unable to complete a request that appears valid. Examples include 500 Internal Server Error, 502 Bad Gateway, 503 Service Unavailable, and 504 Gateway Timeout. In an application-delivery environment, these codes are especially useful for distinguishing problems reported by the application or upstream service from request-side issues. HTTP defines status codes in classes based on their first digit, with the 4xx and 5xx classes reserved for client and server errors, respectively. See the HTTP Semantics status code definitions in RFC 9110 and the MDN HTTP response status reference.

QUESTION NO: 21

Which two must be included in a Wide-IP definition for the Wide-IP to resolve a DNS query? (Choose two.)

A. a name

B. a monitor

C. a load balancing method

D. one or more virtual servers

ANSWER: A C

Explanation:

A Wide IP must have **a name**, because the name is the fully qualified domain name that BIG-IP DNS uses to match an incoming DNS request to the Wide IP object. When a DNS query matches that name, BIG-IP DNS can apply the Wide IP's configuration to select an appropriate destination and construct the DNS response.

A Wide IP also requires **a load balancing method** as part of its resolution behavior. The load-balancing method defines the policy BIG-IP DNS uses when choosing among the available pool members or other eligible resources associated with the Wide IP. For example, the method can favor global availability, round robin selection, or topology-based selection. This policy is fundamental to the Wide IP's DNS-answer selection process; if an administrator does not explicitly select one, the system uses the applicable configured default.

These settings establish the queried DNS identity and the selection logic used to answer matching requests. F5 describes a Wide IP as the DNS name that maps to pools and explains that load-balancing methods determine how BIG-IP DNS selects resources. See the BIG-IP DNS Concepts documentation and the tmsh wideip a reference.

QUESTION NO: 22

Which of the following is NOT a profile type on the BIG-IP?

A. Protocol

B. Application

C. Persistence

D. Authentication

E. SSL

ANSWER: B

Explanation:

Application is the correct choice because it is not a standalone BIG-IP profile type. BIG-IP profiles are configuration objects that attach traffic-handling behavior to a virtual server. The platform organizes these objects into specific profile types and families, including protocol-related profiles, persistence profiles, authentication profiles, and SSL profiles. For example, TCP and UDP profiles define transport-protocol behavior; persistence profiles determine how repeat client requests are directed to the same pool member; authentication profiles define authentication services; and client-side or server-side SSL profiles control TLS processing.

Although BIG-IP can deliver and optimize applications, and it contains profiles for individual application-layer protocols such as HTTP, DNS, FTP, SIP, and others, “Application” by itself is not the name of a selectable profile type. In practice, an administrator chooses the particular profile appropriate to the application protocol and attaches it to the virtual server. This distinction matters because profile selection determines the precise traffic processing feature that BIG-IP applies, rather than merely identifying the workload as an application.

F5’s profile documentation describes profiles as discrete traffic-management configuration objects and documents the available profile categories: BIG-IP LTM Concepts and BIG-IP Local Traffic Management Profiles Reference.

QUESTION NO: 23

Which two of the following statements about how TMOS typically manages traffic between a client and server are accurate? (Choose two.)

A. It changes the destination address before forwarding a request to the server.

B. It changes the destination address before sending a response to the client.

C. It changes the source address before sending a response to the client.

D. It changes the source address before forwarding a request to the server.

ANSWER: A C

Explanation:

TMOS commonly operates as a full proxy, maintaining a client-side connection and a separate server-side connection. When a client sends traffic to a virtual server, the virtual server address is the original destination. To deliver that request to the selected pool member, TMOS creates the server-side flow with the pool member’s address as its destination. Therefore, *It changes the destination address before forwarding a request to the server.* accurately describes the usual destination-address translation from the virtual server to the selected server.

For return traffic, the server-side connection has the pool member as its source. TMOS presents the virtual server address to the client on the client-side flow, rather than exposing the pool member address. Consequently, *It changes the source address before sending a response to the client.* correctly describes the source-address translation used for the response. This behavior lets the BIG-IP system provide a consistent virtual-server endpoint while directing individual requests to pool members. Source NAT can additionally be configured for server-side requests, but it is not the basic address translation that is inherent in the normal virtual-server-to-pool-member and pool-member-to-virtual-server traffic flow. See the F5 virtual server documentation and the F5 pool documentation.

QUESTION NO: 24

The APM Dashboard enables administrators to monitor which two of these metrics? (Choose two.)

A. Number of active sessions

B. Number of new sessions

C. Number of denied users

D. Number of users from each country

ANSWER: A B

Explanation:

The APM Dashboard provides an at-a-glance operational view of Access Policy Manager usage, including session activity trends. **Number of active sessions** is a dashboard metric because it shows the current volume of APM sessions being maintained by the system. This helps administrators assess the present connection load and understand how many users are actively using access services.

Number of new sessions is also monitored on the dashboard. Tracking newly created sessions over time helps reveal authentication and access-demand trends, such as a normal morning login surge or an unexpected increase in connections. Together, active-session and new-session data allow an administrator to distinguish between the current session population and the rate at which users are initiating access.

F5 documents the Access Policy Manager dashboard as a monitoring interface that presents session-related statistics and charts for the selected time interval. These measurements support routine capacity observation and troubleshooting of APM access activity. See the F5 documentation for the Access Policy Manager dashboard and the BIG-IP Access Policy Manager Visual Policy Editor guide.

QUESTION NO: 25

An age-based policy is set up on the ARX that retains only data modified in the last 3 months on tier 1 storage and moves the rest of the data to secondary storage. What happens when the end user tries to access data that has not been touched in 6 months?

A. The end user is unaware that the data has been moved to secondary tier storage and is able to access the data without difficulty.

B. The networking mapping window appears, allowing the end user to re-establish direct access to the data even though it has been moved to secondary storage.

C. An error message appears saying "File is no longer unavailable."

D. A message appears explaining that the file has been archived, and a link to the new secondary storage tier location is provided.

ANSWER: A

Explanation:

The end user is unaware that the data has been moved to secondary tier storage and is able to access the data without difficulty. F5 ARX is a file-virtualization and policy-based data-management platform that presents a consistent namespace to clients while it manages where files physically reside. When the age-based policy relocates a file that has not been modified for six months from tier 1 to secondary storage, ARX retains the file's logical path in the exported namespace. The client therefore continues to request the file by using the same share and pathname, without needing to know the target tier or establish a separate connection to it.

Upon access, ARX transparently directs the request to the file's current physical location on the secondary tier and returns the data through the normal file-access protocol. This transparent access model lets administrators place less-active data on lower-cost storage while preserving application and user access patterns. The migration is consequently a storage-management action rather than a user-visible archive workflow. F5 describes ARX as providing a global namespace and policy-driven file movement across storage resources; see the F5 ARX documentation portal and the F5 ARX data sheet.

QUESTION NO: 26

If a customer has an application that uses a customized protocol, what LTM feature can help optimize the traffic from the application?

A. iRules

B. Network virtual servers

C. HTTP classes

D. Packet filtering

E. Transparent virtual servers

ANSWER: A

Explanation:

iRules can help optimize traffic for an application that uses a customized protocol because they provide programmable, event-driven control over how BIG-IP LTM processes connections and application data. Written in the Tcl language, an iRule can inspect traffic at relevant protocol-processing events and apply tailored logic, such as selecting a pool member, changing persistence behavior, manipulating headers or payload data, redirecting a connection, or enforcing application-specific handling. This flexibility is particularly useful when the protocol is proprietary or otherwise lacks a built-in LTM profile with the necessary optimization and traffic-management behavior. The iRule is attached to the relevant virtual server, allowing the administrator to implement the required custom logic at the traffic-management point without modifying the application itself. F5 documents iRules as a way to customize BIG-IP behavior beyond standard configuration objects and to manage traffic based on events occurring during a connection. See the F5 iRules reference documentation and the iRules introduction for the event-driven model and supported commands.

QUESTION NO: 27

How is persistence configured?

Persistence Table All entries for the one virtual server and pool		
Persistence Values	Member	Age (Seconds)
200.10.0.0	10.10.20.1:80	63
201.12.0.0	10.10.20.3:80	43
153.15.0.0	10.10.20.2:80	76
205.12.0.0	10.10.20.4:80	300
195.64.0.0	10.10.20.3:80	22
198.22.0.0	10.10.20.5:80	176
214.77.0.0	10.10.20.1:80	43

Web_Pool Statistics					
Member	Member Ratio	Member Priority	Outstanding Layer 7 Requests	Connection Count	Status
10.10.20.1:80	3	5	6	18	Available
10.10.20.2:80	3	5	6	12	Available
10.10.20.3:80	3	5	12	5	Disabled
10.10.20.4:80	1	1	8	19	Offline
10.10.20.5:80	1	1	4	9	Available

Virtual Server, Pool and Persistence Profile Settings					
VS_Web_Pool Settings		Web_Pool Settings		Source Persist Settings	
Destination	172.160.22.3:80	Load balancing	Least Connections	Mode	Source Address
Profile(s)	TCP	Priority Activation	Less than 2	Netmask	255.255.0.0
Pool	Web_Pool	Montitir	Done	Timeout	360 seconds
iRules	None				
Persistence	Source_Persist				

A. Persistence is an option within each pool's definition.

B. Persistence is a profile type; an appropriate profile is created and associated with virtual server.

C. Persistence is a global setting; once enabled, load balancing choices are superceded by the persistence method that is specified.

D. Persistence is an option for each pool member. When a pool is defined, each member's definition includes the option for persistence.

ANSWER: B

Explanation:

Persistence is a profile type; an appropriate profile is created and associated with virtual server. On BIG-IP Local Traffic Manager, persistence keeps subsequent connections or requests from a client associated with the same pool member selected for the initial connection. Administrators configure this behavior by creating or selecting a persistence profile, such as source address affinity, cookie, SSL session ID, destination address affinity, or universal persistence, and then associating that profile with the applicable virtual server. The virtual server is the correct attachment point because it defines the client-facing traffic-processing policy, including its pool, protocol profiles, and persistence behavior. A fallback persistence profile can also be configured on the virtual server when appropriate. The selected profile determines how BIG-IP identifies a session and maintains its mapping to a pool member, while the pool continues to provide the available members that can receive traffic. This model lets different virtual servers use distinct persistence methods or timeouts even if they direct traffic to the same pool. F5's documentation describes persistence profiles as the configuration objects used to ensure requests from a client are sent to the same pool member, and documents their assignment in virtual-server configuration. See the F5 persistence profiles reference and F5 persistence configuration documentation.

QUESTION NO: 28

Which statement about root DNS servers is true?

A. Root servers have databases of all registered DNS servers.

B. Root servers have databases of the DNS servers for top-level domains.

C. Root servers have databases of DNS servers for each geographical area. They direct requests to appropriate LDNS servers.

D. Root servers have databases of commonly accessed sites. They also cache entries for additional servers as requests are made.

ANSWER: B

Explanation:

Root servers have databases of the DNS servers for top-level domains. In the DNS hierarchy, root name servers are authoritative for the root zone, which contains delegation information for top-level domains such as `.com`, `.org`, country-code domains such as `.uk`, and infrastructure domains. When a recursive resolver needs to resolve a name and does not already have a usable cached answer, it can query a root server. The root server responds with a referral that identifies the authoritative name servers for the relevant top-level domain. The resolver then queries a top-level domain server, which can refer it to the authoritative servers for the requested domain; those authoritative servers provide the relevant DNS record, such as an address record. Root servers therefore provide the starting point for DNS delegation, rather than maintaining the individual host records for all Internet domains. The root-zone data is coordinated by IANA and is served globally by root-server operators using multiple distributed instances. See the IANA Root Zone Management page and the IANA Root Servers reference for details.

QUESTION NO: 29

Which programming language is the basis for F5 iRules?

A. Lisp

B. C++

C. Java

D. TCL

E. AWK

ANSWER: D

Explanation:

TCL is the basis for F5 iRules. An iRule is a script that BIG-IP executes in response to traffic-processing events, such as a client connection being accepted or an HTTP request being received. F5 implements iRules using the Tool Command Language (Tcl) syntax and extends Tcl with BIG-IP-specific commands, variables, and event contexts. This lets administrators apply custom, event-driven traffic logic directly on the BIG-IP system—for example, inspecting request attributes, selecting a pool member, manipulating headers, or directing traffic according to application requirements. Familiarity with core Tcl constructs, including commands, variables, conditional logic, and string handling, is therefore useful when reading and creating iRules. F5's iRules documentation explicitly describes iRules as being based on Tcl and provides the event and command references used to build this processing logic. See the F5 iRules reference and the BIG-IP iRules concepts documentation for details.

QUESTION NO: 30 - (DRAG DROP)

DRAG DROP

Match these terms with their description.

Select and Place:

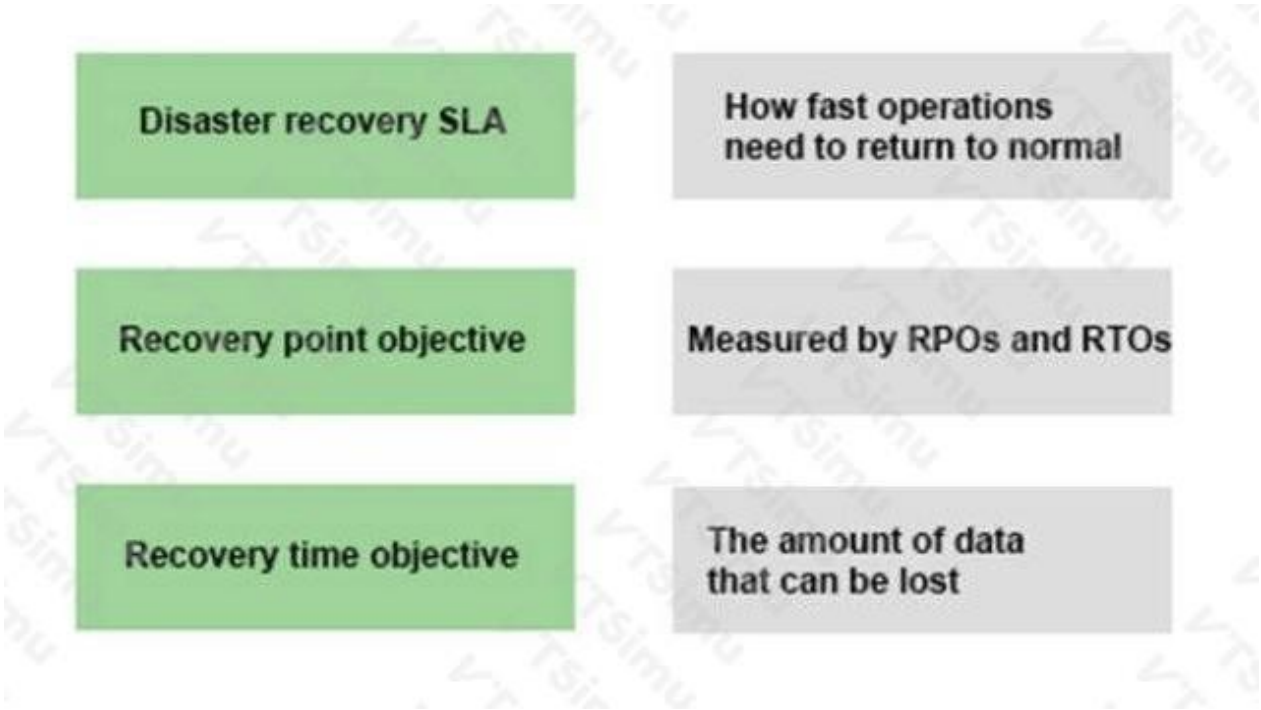

ANSWER:

Explanation:

A disaster recovery service-level agreement defines the recovery commitments an organization expects for an application or service after a disruptive event. Those commitments are normally expressed through two core recovery metrics: the recovery point objective and the recovery time objective. Therefore, “Disaster recovery SLA” matches “Measured by RPOs and RTOs.” This establishes the measurable service expectation for the recovery plan.

The recovery point objective, commonly called RPO, is the maximum acceptable amount of data loss measured in time. For example, an RPO of one hour means that, after recovery, the organization can accept losing up to one hour of recently

created or changed data. This makes “The amount of data that can be lost” the correct description for “Recovery point objective.” The NIST Contingency Planning Guide discusses recovery objectives as key elements of contingency and recovery planning.

The recovery time objective, or RTO, defines how quickly a service must be restored after an interruption. It is concerned with the elapsed downtime that the business can tolerate before normal operations need to resume. Consequently, “How fast operations need to return to normal” correctly matches “Recovery time objective.” Together, RPO addresses the acceptable recovery-data gap and RTO addresses the permitted recovery duration. These measures let a disaster recovery SLA translate business continuity needs into specific, testable operational targets. F5’s disaster recovery glossary provides related continuity and recovery context.

QUESTION NO: 31

Which three statements describe a characteristic of profiles? (Choose three.)

A. Default profiles cannot be created or deleted.

B. Custom profiles are always based on a parent profile.

C. A profile can be a child of one profile and a parent of another.

D. All changes to parent profiles are propagated to their child profiles.

E. While most virtual servers have at least one profile associated with them, it is not required.

ANSWER: A B C

Explanation:

BIG-IP profiles use inheritance to provide reusable traffic-processing configuration. Default profiles are system-supplied baseline objects: administrators use them as the starting point for configuration but do not create or delete those default profile objects. A custom profile is created from a parent profile, allowing it to inherit the parent’s settings and then selectively override settings needed for a particular application or virtual server. This approach avoids repeatedly configuring the same protocol, persistence, SSL, or application-behavior settings.

Profile inheritance can also be layered. A custom profile may inherit from another profile and can itself serve as the parent for additional child profiles. This makes it possible to establish an organization-wide baseline, create an application-specific derivative, and then create further specialized profiles when needed. Inheritance applies to settings that remain inherited; a child setting explicitly overridden locally is maintained at the child level. F5 documents profile inheritance and the use of parent profiles in its BIG-IP Local Traffic Management profile documentation: BIG-IP LTM Profiles Reference and F5 TMSH profile reference.

QUESTION NO: 32

What are Iwo 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. 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 and serial failover paths. “failing over network connection based on heartbeat detection” is correct because a Device Service Cluster uses network failover communication between peer devices. The devices exchange heartbeat and status information over configured failover networks; loss of the required peer communication can cause the standby device to take ownership of traffic groups according to the configured failover behavior. This provides redundancy even when the active unit is no longer available to process application traffic.

“failing over serial cable based on voltage detection” is also correct. Where supported, a directly connected serial failover cable supplies an additional, independent path for determining peer availability. The serial connection detects the electrical state of the peer, including loss of the expected signal/voltage condition, and can contribute to failover decisions. Using both network heartbeat monitoring and a serial failover connection improves resiliency because the devices have more than one mechanism to determine whether the active peer remains operational. F5 documents these failover communication methods and their configuration in its Device Service Clustering guidance: BIG-IP Device Service Clustering Administration and F5 Network Failover Overview.

QUESTION NO: 33

What is the relationship between a BIG-IP node and a pool member?

A. A pool member combines a node address with a service port.

B. A node is a collection of virtual servers.

C. A pool member is a floating self IP address.

D. A node is a DNS load-balancing method.

ANSWER: A

Explanation:

A pool member combines a node address with a service port is correct. A node is an addressable server resource, typically represented by an IP address. A pool member identifies a specific service on that node by associating the node with a port, such as port 80 or port 443.

This distinction allows one node to provide multiple services that can be placed in different pools. F5 explains nodes and pool members in the BIG-IP nodes and pool members documentation.

QUESTION NO: 34

Which three of these scenarios should you identify as being an APM opportunity? (Choose three.)

A. An organization using Novell Netware for authentication.

B. An organization that has been recently fined for failing security compliance.

C. An organization with a traveling sales force.

D. An organization with a single location and no remote employees.

E. An organization that needs to ensure users are using Symantec antivirus software.

F. An organization sharing a public Web site for all Internet users.

ANSWER: B C E

Explanation:

An organization that has been recently fined for failing security compliance, an organization with a traveling sales force, and an organization that needs to ensure users are using Symantec antivirus software are strong Access Policy Manager (APM) opportunities. APM provides centralized, context-aware access control for applications, so it is well suited to organizations

that must demonstrate and enforce security policy consistently. Compliance-driven organizations can use access policies to apply authentication, authorization, logging, and session controls before allowing access to protected resources.

A traveling sales force creates a clear need for secure remote access to corporate applications from varied locations and devices. APM can deliver identity-aware remote access and apply the same access rules regardless of where the user connects. Endpoint security requirements are also directly aligned with APM capabilities: endpoint inspection and posture assessment can evaluate client-device characteristics, including antivirus-related checks, and use the result in an access decision. This lets the organization require an acceptable security posture before granting access. F5 describes APM as providing secure, context-aware access to applications, while its policy framework supports authentication and endpoint-related access decisions. See F5 Access Policy Manager and F5 APM Visual Policy Editor documentation.

QUESTION NO: 35

When a business is hacked, they often lose more than money. What are the other consequences to a business as a result of being hacked? Select two.

A. Helpful third party reviews of the security needs of the customer's web applications.

B. Valuable free press that occurs as companies address hacking incidents.

C. Penalties related to non-compliance with laws and regulations.

D. Final resolution of all security vulnerabilities of the business' web applications.

E. Loss of customers when slow connections drive customers to competitor's site.

ANSWER: C E

Explanation:

Penalties related to non-compliance with laws and regulations are a significant consequence of a successful hack. Organizations that fail to protect regulated information, such as personal, payment, or health data, can face investigations, mandatory corrective actions, contractual consequences, and regulatory fines. The applicable requirements vary by jurisdiction and industry, but breach-response obligations commonly include notification, documentation, and demonstrated security remediation. The U.S. Federal Trade Commission describes the importance of implementing reasonable safeguards and responding appropriately to data-security incidents: FTC Data Security guidance.

Loss of customers when slow connections drive customers to competitor's site is also a genuine business impact. An attack can exhaust application, network, or infrastructure capacity and degrade availability and response time. Even when an application remains technically reachable, a poor user experience can interrupt purchases, reduce engagement, damage customer confidence, and cause users to choose a competitor. F5 explains that distributed denial-of-service attacks can affect the availability and performance of online services, making resilience and traffic protection essential to application delivery: F5: Distributed Denial of Service (DDoS). These legal and customer-retention impacts can persist beyond the immediate technical incident.

QUESTION NO: 36

Which persistence method will always recognize a client when the client returns to the same virtual server?

A. SSL

B. MSRDP

C. expression [universal]

D. No persistence method work in all situations.

E. source address

ANSWER: D

Explanation:

No persistence method work in all situations. BIG-IP persistence is based on information available in a particular client connection or application transaction, together with a persistence record maintained by the system. That information is not guaranteed to remain available, unique, or unchanged whenever a client returns. For example, a client can reconnect through a different network path, an intermediary can alter or conceal identifying information, or the application can start a new session without presenting the value used to associate it with an earlier connection. In addition, persistence records have lifetimes and can be removed when their configured timeout expires or when relevant configuration or traffic conditions change.

Consequently, persistence must be selected according to the application protocol, the available client-identifying data, and the required session behavior; it is not an absolute guarantee of recognition in every possible return scenario. F5 describes persistence as a mechanism for directing subsequent client requests to the same pool member based on configured persistence criteria and records, rather than as universal client identification. See the F5 BIG-IP Local Traffic Management persistence documentation and F5 guidance on persistence profiles.

QUESTION NO: 37

When an administrator creates a new access policy in the Visual Policy Editor, which three options are included by default? (Choose three.)

A. A fallback option

B. An Allow Ad box

C. A Deny End box

D. An empty Resource Assign item

E. A Start box

F. A Block All option

ANSWER: A C E

Explanation:

A newly created access policy in the BIG-IP Access Policy Manager Visual Policy Editor begins with a minimal, secure policy flow. A Start box is included as the policy entry point. It represents where policy evaluation begins for a client request and is the location from which an administrator adds logon, authentication, endpoint inspection, assignment, and other access-policy actions.

A fallback option is also created by default. The fallback branch provides the default path through the visual policy when no explicitly configured branch expression is matched. This gives the policy a defined continuation path while the administrator builds and connects the required access decisions.

A Deny End box is included as the default terminal result. This secure default ensures that a session is not granted access merely because a policy has been created; access is only allowed after the administrator adds the appropriate checks and connects a successful path to an allow result. Together, the Start box, fallback option, and Deny End box form the initial policy skeleton that can then be expanded in the Visual Policy Editor. F5 documents Visual Policy Editor policy construction and policy endings in the BIG-IP Access Policy Manager Visual Policy Editor documentation and the F5 Access Policy Manager knowledge documentation.

QUESTION NO: 38

Monitors can be assigned to which three resources? (Choose three.)

A. NATs

B. pools

C. iRules

D. nodes

E. SNATs

F. pool members

G. virtual servers

ANSWER: B D F

Explanation:

On BIG-IP, health monitors are used to determine the availability of application resources and can be associated at several levels of the load-balancing hierarchy. **pools** can have monitors assigned so that BIG-IP evaluates the health of the pool's members according to the configured monitoring method. This is useful when the same health-check policy should apply to the resources participating in a pool.

nodes can also have monitors assigned. A node represents a server address, and node monitoring provides a way to determine whether that server is generally reachable and available, independently of a particular service port.

pool members support monitor assignment as well. Because a pool member combines a node address with a service port, a monitor at this level can verify the availability of the specific application service that receives traffic. BIG-IP uses the resulting monitor status when making load-balancing decisions, directing new connections only to members considered available. Monitor assignment and inheritance allow health checks to be applied at the level that best matches the required availability scope. See the BIG-IP monitor reference and BIG-IP pool and member documentation.

QUESTION NO: 39

The main drawback to using an APM Deployment Wizard is that it is not possible to edit the resulting objects, including the access policy.

A. True

B. False

ANSWER: B

Explanation:

False is correct because an APM Deployment Wizard creates a configured access-policy deployment that administrators can subsequently inspect, modify, and extend. After completing a wizard, the resulting access profile is available under *Access Policy* > *Access Profiles*, where its visual policy can be opened in the Visual Policy Editor. Administrators can add or change authentication and authorization actions, configure endpoint checks, alter logon-page settings, edit resource assignments, and adjust other policy behavior to fit requirements that were not captured by the initial wizard selections.

The wizard is intended to simplify and accelerate creation of common APM configurations, rather than permanently lock the generated objects. Its output provides a starting configuration, while the normal APM management interfaces remain available for ongoing administration. F5's APM documentation describes editing an access policy by opening the applicable access profile and selecting *Edit* for the access policy, which confirms that policies created through standard configuration workflows can be maintained after deployment. See the F5 APM Visual Policy Editor documentation and the F5 APM access-policy reference.

QUESTION NO: 40

Which event is always triggered when the client sends data to a virtual server using TCP?

A. HTTP_DATA

B. CLIENT_DATA

C. HTTP_REQUEST

D. VS_CONNECTED

ANSWER: B

Explanation:

`CLIENT_DATA` is the correct event because it is raised when BIG-IP receives application data from the client side of a TCP connection. It is a general TCP payload event and is not dependent on the payload being formatted as a particular application protocol. This makes it the appropriate iRules event for inspecting, collecting, modifying, or otherwise acting on client-to-virtual-server data at the TCP layer.

The event can occur more than once during a connection because TCP is a byte-stream protocol: application content can arrive in multiple segments, and BIG-IP processes data as it becomes available. An iRule commonly uses commands such as `TCP::collect` to control how much client data is buffered before handling it in `CLIENT_DATA`. This behavior allows an iRule to operate on arbitrary TCP services, including custom protocols as well as protocols carried over TCP.

F5 documents `CLIENT_DATA` as occurring when the system receives data from the client, and its iRules documentation describes the TCP data-collection mechanisms used with this event. See the F5 CLIENT_DATA event reference and the F5 TCP::collect reference.

QUESTION NO: 41

Which two can be a part of a virtual server's definition? (Choose two.)

A. rule(s)

B. pool(s)

C. monitor(s)

D. node address(es)

E. load balancing method(s)

ANSWER: A B

Explanation:

A virtual server can include **rule(s)**, meaning iRules associated with that virtual server. An iRule is a Traffic Management Microkernel script that is attached to a virtual server and executes in response to relevant traffic events. This lets the virtual server apply custom traffic-processing logic, such as selecting traffic behavior based on connection or HTTP request properties.

A virtual server can also include **pool(s)**. In the usual configuration, the virtual server specifies a default pool to which it sends eligible client connections. A pool represents a collection of pool members that can receive application traffic. The BIG-IP system uses the pool configuration to select an available member for each connection, while additional traffic logic can select among pools where required. Thus, both iRules and pools are configurable objects that may be associated with a virtual server definition.

F5 documents virtual-server properties, including the default pool and associated iRules, in its configuration documentation. See the BIG-IP virtual server tmsh reference and the F5 BIG-IP documentation portal.

QUESTION NO: 42

Which three properties can be assigned to nodes? (Choose three.)

A. ratio values

B. priority values

C. health monitors

D. connection limits

E. load balancing mode

ANSWER: A C D

Explanation:

Nodes represent the individual servers or server addresses that a BIG-IP system can use to deliver application traffic. The node-level properties in this question are **ratio values**, **health monitors**, and **connection limits**. A ratio value is an administrative weighting attribute that can be used when a pool's load-balancing method considers ratios, allowing the BIG-IP system to distribute more connections to a node with a higher relative ratio. Health monitors can be associated with a node so that BIG-IP can actively determine whether that server address is available; the resulting monitor status contributes to whether traffic should be sent to the node. A connection limit can also be configured for a node, establishing the maximum number of concurrent connections the BIG-IP system should allow to that node. These settings are node attributes, although their practical effect can depend on the pool and load-balancing configuration in which the node is used. F5's tmsh node reference documents the `ratio`, `monitor`, and `connection-limit` properties: F5 BIG-IP tmsh ltm node reference. For background on configuring nodes and pools, see F5 Local Traffic Manager node documentation.

QUESTION NO: 43

LTM runs on_____F5's proprietary product platform. (Fill in the correct answer)

A. ARX

B. Firepass

C. Acopia

D. TMOS

ANSWER: D

Explanation:

TMOS is correct because BIG-IP Local Traffic Manager (LTM) is an application delivery and load-balancing module that operates on F5's Traffic Management Operating System (TMOS) platform. TMOS provides the common operating environment and shared services used by BIG-IP modules, including the full-proxy traffic architecture, networking functions, security capabilities, high availability, and the management plane. LTM uses these TMOS capabilities to inspect application flows, make traffic-distribution decisions, apply persistence and health monitoring, and direct clients to available pool members. In practical administration, configuration objects such as virtual servers, pools, nodes, monitors, profiles, and traffic policies are created and managed within the BIG-IP TMOS system. This relationship is fundamental to F5 terminology: LTM identifies the local traffic-management functionality, while TMOS identifies the proprietary software platform on which that functionality is delivered. F5 describes LTM as part of its BIG-IP application delivery offerings and identifies TMOS as the operating system supporting BIG-IP services. See F5's Local Traffic Manager overview and its BIG-IP LTM technical documentation for additional product and platform context.

QUESTION NO: 44

Which two capabilities result from the BIG-IP full-proxy architecture? (Choose two.)

A. Independent client-side and server-side TCP connections

B. Independent traffic optimization on each side of the proxy

C. A single TCP connection from client to server

D. Identical TCP profiles on both sides of every connection

E. Direct packet forwarding without connection state

ANSWER: A B

Explanation:

Independent client-side and server-side TCP connections and **independent traffic optimization on each side of the proxy** are correct. BIG-IP terminates the client-side connection and establishes a separate connection to the selected server. Because the connections are separate, BIG-IP can apply different TCP behavior and optimization appropriate to the client network and the server network.

BIG-IP does not simply forward a single end-to-end TCP session, and it does not require the client and server to negotiate identical TCP settings. F5 describes this architecture in the BIG-IP platform training material.

QUESTION NO: 45

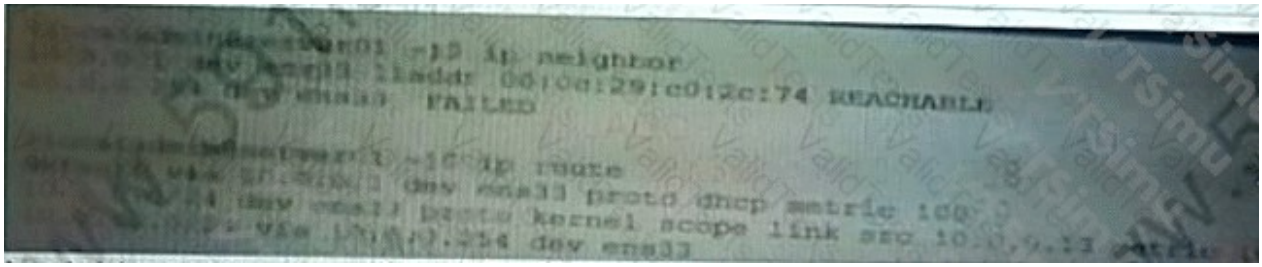

An administrator is trying to debug why the 10.0.10.0/24 network is inaccessible from a certain server. The server is able to access the internet Successfully.

The administrator inspects the routing table and ARP table on the affected server.

Given the terminal output shown, what could be causing this issue?

A. The default gateway is not sending ARP responses

B. The static route for 10. 0.10.0/24 is misconfigured

C. The local server is unable to receive ARP responses

D. The default route (0. 0. 0. 0/0) is misconfigured

ANSWER: B

Explanation:

The static route for 10. 0.10.0/24 is misconfigured is the supported conclusion. Because the server can access the internet, its default route and the Layer 2 resolution needed to reach the default gateway are functioning sufficiently for off-network traffic. Traffic destined specifically for 10.0.10.0/24 is instead evaluated against the more-specific route for that destination prefix. A route for a /24 network takes precedence over the default route because routers use longest-prefix matching when selecting a forwarding path. Therefore, an incorrect next-hop address, incorrect outgoing interface, or otherwise unusable configuration on that static route can prevent access to 10.0.10.0/24 while leaving general internet access unaffected. The routing-table and ARP information together help isolate the problem to the destination-specific forwarding entry rather than the server's normal default-gateway path. F5 BIG-IP follows the same fundamental routing behavior: it uses the most specific matching route when multiple routes can match a destination. See the F5 documentation on route configuration and route selection and the Linux ip-route documentation for route lookup behavior.

QUESTION NO: 46

What are some changes that must be made on the GTM System so that log messages are sent to centralized System Log servers?

A. The IP address of the server must be added to the wideip.conf file.

B. The IP address of the server must be added to the syslog-ng.conf file.

C. The IP address of the server and valid user id/password combination must be added to the hosts.allow file.

D. The IP address of the server and valid user id/password combination must be added to the syslog ng.conf file.

ANSWER: B

Explanation:

The IP address of the server must be added to the syslog-ng.conf file. GTM (now known as BIG-IP DNS) uses the syslog-ng logging service to process local system messages and forward selected messages to remote, centralized syslog collectors. A remote syslog destination is defined in the syslog-ng configuration, where the collector's address, transport protocol, and (where applicable) port are specified. Once the remote destination is configured and associated with the relevant log source or filter, GTM system log messages can be transmitted to that centralized server for retention, monitoring, and correlation.

In supported BIG-IP versions, administrators should generally make this configuration through the BIG-IP configuration interfaces, such as the `sys syslog` tmsh object or the corresponding Configuration utility settings, rather than manually editing generated system files. Those interfaces maintain the underlying syslog-ng configuration and help ensure that the setting persists correctly across service reloads and configuration changes. F5 documents the remote-server settings, including the server address and transport details, in the tmsh sys syslog reference. The BIG-IP DNS product documentation is available through the F5 BIG-IP DNS documentation portal.

QUESTION NO: 47

Which of the following does NOT describe network performance management?

A. Updating software on various devices

B. Seeing graphs in near real-time

C. Visibility into multiple network devices

D. Being alerted to an issue as it happens

E. Viewing virtual server statistics

ANSWER: A

Explanation:

Updating software on various devices is the activity that does not describe network performance management. Performance management is concerned with monitoring and evaluating how the network and its application-delivery components are operating. It uses collected statistics, health indicators, traffic data, and event information to establish normal behavior, identify degradation, and support troubleshooting. For example, an administrator may monitor throughput, connection rates, latency-related indicators, interface utilization, resource consumption, and virtual-server statistics to assess the performance and availability of services.

Software updating, by contrast, is an administrative maintenance and lifecycle-management task. Although keeping software current can improve reliability, security, and feature support, the act of deploying or updating software is not itself the ongoing observation, measurement, alerting, and analysis represented by network performance management. F5 BIG-IP provides statistics and monitoring data for objects such as virtual servers, pools, and pool members, enabling operators to evaluate application delivery behavior and respond to operational conditions. See the F5 BIG-IP virtual-server statistics reference and the F5 BIG-IP monitoring documentation.

QUESTION NO: 48

Which function does a caching profile provide on the BIG IP server?

A. caches client IP addresses to improve response times

B. caches HTTP headers lo improve server load times.

C. caches selected pool members to improve load balancing times.

D. caches data for reuse to cut down on transfer time for future requests

ANSWER: D

Explanation:

A caching profile provides the ability to store cacheable application content and serve that stored content in response to later matching client requests. On BIG-IP, the Web Acceleration/Intelligent Traffic Management Cache (RAM Cache) functionality keeps objects in memory after they are retrieved from an origin server or pool member. When another request can be satisfied by a valid cached object, BIG-IP returns the object directly rather than requesting it from the pool member again. This reduces repeated application-server processing and avoids unnecessary transfers between BIG-IP and the origin infrastructure, which can improve response time for clients and reduce server load. Cache behavior is governed by the caching profile together with HTTP caching information, such as response cache-control directives and object expiration. The cached object is therefore reusable content, rather than client-address information, load-balancing selections, or merely a collection of headers. F5 describes RAM Cache as an in-memory cache that can cache HTTP content and provide that content directly to clients when appropriate. See the F5 RAM Cache profile documentation and the BIG-IP Local Traffic Management Profiles Reference.

QUESTION NO: 49

Which two of the following are costs businesses may face in dealing with unstructured data? (Choose two.)

A. Lost productivity due to server downtime

B. Buying backup media

C. Buying additional storage capacity

D. Paying to convert unstructured data into structured data

ANSWER: B C

Explanation:

Unstructured data, such as documents, images, video, email, and log files, can grow rapidly because it does not fit neatly into traditional database fields and is often retained for operational, legal, analytical, or business reasons. Consequently, **Buying additional storage capacity** is a direct cost: organizations need sufficient primary storage to hold expanding data sets, as well as capacity that supports performance, availability, and future growth. IBM describes unstructured data as information that lacks a predefined data model and commonly includes media files and text-heavy content, which can create substantial data-management requirements.

Buying backup media is also a direct cost because valuable unstructured data must be protected against accidental deletion, corruption, hardware failures, ransomware, and site-level events. Backup processes may require storage media, backup repositories, replication capacity, retention management, and associated infrastructure. As the volume of source data increases, the amount of protected backup capacity generally increases as well. These costs are central to planning a resilient data-protection strategy, where organizations must maintain recoverable copies while managing retention and storage consumption. See IBM: What is unstructured data? and NIST resources on enterprise data and system-management resilience.

QUESTION NO: 50

You need to terminate client SSL traffic at the BIG-IP and also to persist client traffic to the same pool member based on a BIG-IP supplied cookie. Which four are profiles that would normally be included in the virtual server's definition? (Choose four.)

A. TCP

B. HTTP

C. HTTPS

D. ClientSSL

E. ServerSSL

F. CookieBased Persistence

ANSWER: A B D F

Explanation:

TCP, HTTP, ClientSSL, and CookieBased Persistence are the normal profile set for this virtual-server requirement. A TCP profile provides the Layer 4 connection handling used by a standard TCP virtual server. The HTTP profile enables BIG-IP to parse and process HTTP traffic after the encrypted client connection has been decrypted. This HTTP awareness is required for BIG-IP to insert, read, and act on a persistence cookie.

The ClientSSL profile is attached to the client side of the virtual server and contains the certificate, private key, and TLS settings that allow BIG-IP to establish the client-facing TLS session and decrypt the incoming HTTPS request. CookieBased Persistence supplies a BIG-IP-managed cookie persistence method, allowing BIG-IP to associate subsequent requests carrying that cookie with the previously selected pool member. Together, these profiles let BIG-IP accept TLS from clients, inspect the resulting HTTP requests, and maintain a consistent server selection for a client session. F5 documents Client SSL configuration in its Client SSL profile documentation and describes cookie persistence configuration in the Cookie persistence documentation.

QUESTION NO: 51

When initially configuring the BIG-IP system using the config utility, which two parameters can be set. (Choose two.)

A. the netmask of the SCCP

B. the IP address of the SCCP

C. the port lockdown settings for the SCCP

D. the netmask of the host via the management port

E. the IP address of the host via the management port

F. the port lockdown settings for the host via the management port

ANSWER: D E

Explanation:

The initial BIG-IP configuration utility establishes the device's out-of-band management connectivity. The IP address of the host via the management port is a required foundational setting because it gives administrators a reachable address for subsequent administrative access, including HTTPS access to the Configuration utility and SSH access where enabled. The netmask of the host via the management port is configured with that address so the BIG-IP system can determine which management destinations are on its directly connected network. Together, these values define the Layer 3 configuration for the dedicated management interface.

This initial setup workflow is deliberately focused on basic platform and management-network settings, allowing the administrator to securely reach the system before configuring application traffic objects such as VLANs, self IP addresses, virtual servers, pools, and traffic-management policies. F5's setup documentation describes assigning the management interface address and network mask as part of the initial system configuration. See the F5 BIG-IP Virtual Edition setup documentation and F5 guidance for configuring a BIG-IP management IP address.

QUESTION NO: 52

Which three statements are true about SNATs? (Choose three.)

A. SNATs provide bidirectional traffic initiation.

B. SNATs support UDP, TCP, and ICMP traffic.

C. SNATs provide a many-to-one mapping between IP addresses.

D. SNAT addresses can be identical to virtual server IP addresses.

ANSWER: A B C

Explanation:

SNATs provide bidirectional traffic initiation because the BIG-IP system maintains translation and connection state for traffic passing through the device. This enables translated flows to be handled in both directions, including the return traffic associated with a connection. SNATs support common IP traffic protocols, including TCP, UDP, and ICMP, so source-address translation can be applied to these traffic types according to the configured SNAT policy or object. SNAT also provides many-to-one address mapping: multiple original client source addresses can be translated to one shared source address. When port translation is used, the BIG-IP system differentiates simultaneous flows by using protocol and port information in its connection table. This allows private or otherwise non-routable client addresses to communicate with destination servers while presenting a routable translated source address. F5 documents SNAT as a source-address translation feature and describes its use for translating client addresses and managing associated traffic flows. See the F5 BIG-IP SNAT configuration documentation and the F5 SNAT API reference.

QUESTION NO: 53

Which two of the following LTM load balancing methods require the least amount of resources? (Choose two.)

A. Round robin

B. Ratio

C. Observed

D. Fastest

E. Predictive

F. Least connections

ANSWER: A B

Explanation:

Round robin and Ratio require the least system resources because they use static selection logic rather than continuously gathering and evaluating changing performance data. Round robin cycles through the available pool members in sequence, maintaining only the information needed to identify the next member. Ratio uses administrator-configured relative weights to distribute new connections proportionally among eligible members. Both methods therefore make a straightforward selection at connection time without needing to track each member's current connection count, application response behavior, or historical performance trends.

These methods are appropriate when pool members have similar capacity, or when their relative capacity can be represented with configured ratios. The BIG-IP pool configuration documentation lists round-robin and ratio as supported load-balancing modes and describes ratio as distributing connections according to configured member ratios. F5's load-balancing-method documentation also distinguishes these static approaches from dynamic methods that base decisions on measured runtime conditions. See F5 BIG-IP tmsh ltm pool reference and F5 load balancing methods documentation.

QUESTION NO: 54

What is NOT a benefit of using a SNAT?

A. ASM can be deployed easily

B. No changes are needed on the servers

C. Fail open is easy to add

D. Higher performance than other configuration

ANSWER: D

Explanation:

Higher performance than other configuration is not an inherent benefit of Source Network Address Translation (SNAT). SNAT changes the source address, and usually the source port, of a connection as BIG-IP forwards traffic to the destination. BIG-IP must maintain translation and connection state so that return traffic can be mapped back to the original client correctly. This is a functional networking capability intended to solve routing and addressability requirements, rather than a mechanism designed to improve throughput or connection-processing performance over other valid deployment designs.

BIG-IP platforms are engineered to process SNAT traffic efficiently, and actual performance depends on the appliance or virtual edition resources, enabled services, traffic profile, connection rate, SSL processing, security inspection, and the broader network design. Therefore, SNAT may be appropriate and perform well in a particular architecture, but it cannot accurately be presented as providing higher performance than other configurations as a general benefit. F5's SNAT documentation describes the feature as source-address translation used to ensure return traffic flows through BIG-IP and to support network-routing designs. See F5 KB K7820: Overview of SNAT features and F5 CloudDocs: SNAT iRule command.

QUESTION NO: 55

Assume the bigd daemon fails on the active system. Which three are possible results? (Choose three.)

A. The active system will restart the bigd daemon and continue in active mode.

B. The active system will restart the TMM daemon and continue in active mode.

C. The active system will reboot and the standby system will go into active mode.

D. The active system will failover and the standby system will go into active mode.

E. The active system will continue in active mode but gather member and node state information from the standby system.

ANSWER: A C D

Explanation:

The possible outcomes are that the active system restarts the `bigd` daemon and remains active, that the active system reboots and its peer becomes active, or that the active system fails over and the standby system becomes active. The `bigd` process performs the health-monitor probing used to determine pool-member and node availability. BIG-IP system services supervise critical daemons and can attempt to restore a failed process. Consequently, a recoverable `bigd` failure can result in a daemon restart without changing the device's active traffic-group role.

If the failure is considered persistent or meets the configured high-availability failure criteria, BIG-IP can protect application availability by relinquishing the active role to the peer. Depending on the configured daemon high-availability action and the severity of the event, this protection can be a failover or a reboot of the affected active device; in either case, the peer can transition to active and take ownership of the traffic group. This behavior aligns health-monitor operation with the broader device-service-clustering model, which is designed to preserve service continuity when an active unit experiences critical service failures. See F5's BIG-IP monitoring reference and Device Service Clustering administration documentation.

QUESTION NO: 56

Learning suggestions in the Policy Building pages allow for which of the following? (Choose two.)

A. XML-based parameters and associated schema are automatically learned.

B. Blocking response pages can be automatically generated from web site content.

C. URL-level parameters are displayed when found and can be accepted into the current policy.

D. The administrator may modify whether the BIG-IP ASM System will learn, alarm, or block detected violations.

E. Maximum acceptable values for length violations are calculated and can be accepted into the security policy by the administrator.

ANSWER: C E

Explanation:

Learning suggestions in BIG-IP Application Security Manager's Policy Building workflow are derived from observed application traffic and let an administrator incorporate discovered application characteristics into the security policy. **URL-level parameters are displayed when found and can be accepted into the current policy.** ASM learns parameters associated with application URLs and presents them as actionable suggestions, so the policy can be updated to recognize expected request structure without manually defining every parameter in advance.

Maximum acceptable values for length violations are calculated and can be accepted into the security policy by the administrator. As ASM observes legitimate traffic, it can develop suggested parameter length limits based on the values it sees. The administrator reviews and accepts the recommendation, which adds an appropriate parameter-length constraint to the policy. This supports positive security by allowing the policy to reflect the application's normal input patterns while keeping enforcement under administrative control. F5's documentation describes traffic learning as a process that produces policy suggestions from observed requests and allows those suggestions to be reviewed and accepted into the policy. See F5: Learning about traffic and F5: Using the automatic learning feature.

QUESTION NO: 57

Which two of the following statements are accurate descriptions of the ARX index? (Choose two.)

A. The ARX index stores the bulk of its data as bit arrays and answers most questions by performing bitwise logical operations on these bitmaps.

B. The ARX index contains the index keys in sorted order, with the leaf level of the index containing the pointer to the page and the row number in the data page.

C. The ARX index is the key enabler of transparent data mobility because the ARX tracks where files reside at all times, from creation through migration and deletion.

D. The ARX index is 100% disposable and can be rebuilt at any point in time without disrupting client or application access to data.

ANSWER: C D

Explanation:

The ARX index is the metadata foundation that enables the ARX system to present a stable, global file namespace while files can be placed, migrated, or removed across managed storage. "The ARX index is the key enabler of transparent data mobility because the ARX tracks where files reside at all times, from creation through migration and deletion" is therefore accurate. ARX uses this continuously maintained file-location knowledge to direct client requests to the appropriate physical file location while allowing administrators to apply lifecycle and tiering policies without changing the path used by clients.

"The ARX index is 100% disposable and can be rebuilt at any point in time without disrupting client or application access to data" is also accurate. The index is derived metadata rather than the authoritative content store; ARX can reconstruct it by rediscovering the exported file systems and their contents. This design supports resilience and recoverability while preserving access through the ARX namespace. F5's ARX materials describe the platform's file virtualization and transparent movement capabilities, and F5's documentation portal provides product documentation resources: F5 ARX Data Sheet and F5 Technical Documentation.

QUESTION NO: 58

An administrator needs to load balance connections to the application server with the lowest number of

Which load balancing method meets these requirements?

A. Least Sessions

B. Ratio Least Connections (member)

C. Least Connections (member)

D. Least Connections (node)

ANSWER: D

Explanation:

Least Connections (node) is the appropriate method when the requirement is to select the application server, represented by an F5 node, that currently has the fewest active connections. This method evaluates the connection count at the node level rather than limiting its calculation to a particular pool member. That distinction is important when the same application server is used by more than one pool or has multiple pool-member entries: the BIG-IP system can account for the node's overall active connection load when making the load-balancing decision.

F5 defines the Least Connections (node) method as selecting the node with the fewest active connections. It is therefore suited to distributing new client connections toward the least-busy application server, helping avoid concentrating connections on a server that is already handling more traffic. The connection count used for this decision is dynamic, so the selected node can change as connections are established and closed. F5's load-balancing-method reference describes the behavior and scope of node-level and member-level methods: F5 KB K9573: Overview of BIG-IP load balancing methods. The available pool load-balancing modes are also documented in the F5 BIG-IP tmsh ltm pool reference.

QUESTION NO: 59

A BIG-IP Administrator is cleaning up unused Virtual Servers and pools via the Configuration Utility while trying to delete certain administrator receives an error message that indicates that the pool is in use.

What is the likely cause of this error?

A. The pool members are all disabled

B. The pool is in use by a Virtual Server

C. The pool members are marked up with a green circle

D. The pool is in use by a monitor

ANSWER: B

Explanation:

The pool is in use by a Virtual Server is correct because a BIG-IP pool can be assigned as the default pool for a virtual server. That virtual server configuration creates a direct object dependency: the virtual server refers to the pool whenever it processes matching client connections and selects a pool member for load balancing. BIG-IP prevents deletion of configuration objects that are still referenced, preserving configuration integrity and preventing the virtual server from pointing to an object that no longer exists. Before deleting the pool, the administrator should identify every virtual server that references it, remove or replace the pool association as appropriate, and then save the configuration. In the Configuration utility, virtual-server settings expose the configured default pool, making this association straightforward to review. F5's documentation describes that a virtual server can specify a default pool and that the system load balances traffic to members of that pool. See F5 BIG-IP virtual server documentation and F5 BIG-IP pool documentation.

QUESTION NO: 60

Select the key reasons F5 is able to handle DNS DDoS attacks so effectively? Select two.

A. F5 can ensure a DNS DDoS attack is not successful.

B. F5 has high performance DNS services.

C. F5 can answer the DNS queries directly.

D. With Global Traffic Manager (GTM), F5 completely stops all DNS DDoS attacks.

E. F5 can ensure a customer never faces a DNS DDoS attack.

ANSWER: B C

Explanation:

F5's effectiveness against DNS-focused DDoS events is based on its ability to provide high-capacity DNS processing and to respond to requests at the DNS tier rather than allowing every query to reach the authoritative DNS infrastructure. "F5 has high performance DNS services" is correct because BIG-IP DNS is designed to process DNS traffic at scale, helping organizations absorb exceptionally high query rates and maintain DNS availability during volumetric query floods. This performance capacity is fundamental to mitigating attacks intended to exhaust DNS server resources or network capacity.

"F5 can answer the DNS queries directly" is also correct. BIG-IP DNS can use DNS Express to maintain a rapid, synchronized copy of DNS zone data and answer authoritative DNS requests locally. Serving valid responses from this high-performance DNS layer reduces dependency on the original authoritative servers, helping protect those servers from overload while legitimate users continue to receive DNS answers. Together, scalable query handling and direct DNS responses provide resilience and offload during DNS DDoS conditions. F5 describes its DNS services and DNS load-balancing capabilities at F5 DNS Services and documents BIG-IP DNS concepts at BIG-IP DNS Services Implementations.

QUESTION NO: 61

in which scenario is a full proxy TCP connection required?

A. when SSL offloading is used

B. when Traffic is routed

C. when administrators manage the servers directly

D. when 3 server responds directly to the client

ANSWER: A

Explanation:

when SSL offloading is used is correct because SSL offloading requires the BIG-IP system to terminate the client-side TLS/SSL session itself. The device must therefore act as a full proxy: it accepts and maintains a TCP connection from the client, performs the TLS handshake and decryption using the configured client SSL profile, and creates a separate connection toward the selected pool member. Depending on the virtual server configuration, the server-side connection can be sent in clear text or protected through a separate server SSL profile. This separation enables BIG-IP to inspect decrypted HTTP traffic, apply traffic-management policies, select a pool member, reuse or independently manage server-side connections, and then encrypt responses as needed for the client connection. In contrast to simple packet forwarding, the client and server do not share one end-to-end TCP session; BIG-IP is the TCP endpoint for each side. F5 describes BIG-IP LTM as a full-proxy architecture that independently manages client-side and server-side flows, which is fundamental to SSL/TLS termination and re-encryption. See the F5 BIG-IP LTM Concepts documentation and F5's full-proxy architecture overview.

QUESTION NO: 62

What describes the third â€˜â€™Aâ€™â€™ in the common authentication acronym AAA?

A. provides redundancy

B. measures usage against an identity

C. provides user filtered access

D. ensures the correct identity

ANSWER: B

Explanation:

measures usage against an identity is correct because the third term in AAA is *Accounting*. Accounting records actions and resource consumption associated with an authenticated identity. Depending on the system and service, these records can include login and logout times, session duration, commands issued, data transferred, or other auditable usage details. This information supports operational reporting, auditing, capacity planning, chargeback, and security investigations by connecting activity to a specific user or account. In an application-delivery environment, AAA services are commonly integrated with centralized identity systems so that access activity can be logged consistently across protected applications and services. The Internet Engineering Task Force describes accounting as collecting and recording resource-consumption information, including the identity associated with the usage. F5 Access Policy Manager documentation likewise discusses using authentication and authorization policies with AAA infrastructure to manage and track user access. See IETF RFC 2903: Generic AAA Architecture and F5 BIG-IP Access Policy Manager documentation.

QUESTION NO: 63

Which of the following are properties of an ASM logging profile? (Choose three.)

A. storage type

B. storage filter

C. storage policy

D. web application

ANSWER: A B C

Explanation:

An ASM logging profile controls how Application Security Manager event data is selected and retained when the system records traffic and security events for an attached security policy. **storage type** is a profile property because it defines the category and handling of stored logging data. **storage filter** is a key property because it determines which requests and security events meet the criteria to be logged, allowing administrators to focus storage and analysis on meaningful traffic. **storage policy** is also a property because it governs the policy applied to retained logging information, including how stored data is managed over time. Together, these settings let an administrator balance forensic visibility, reporting requirements, and available system resources while preserving the events needed to investigate application attacks. The logging profile is then associated with the relevant ASM-protected traffic configuration so that the selected logging behavior is applied consistently. F5's ASM documentation describes configuring logging profiles to control the storage and filtering of Application Security events, while the BIG-IP ASM implementation documentation provides the related configuration workflow. See F5 BIG-IP ASM: Logging Application Security Events and F5 Support: Configuring logging profiles.

QUESTION NO: 64

Assume a virtual server has a ServerSSL profile. What SSL certificates are required on the BIG-IP?

A. No SSL certificates are required on the BIG-IP.

B. The BIG-IP's SSL certificates must only exist.

C. The BIG-IP's SSL certificates must be issued from a certificate authority.

D. The BIG-IP's SSL certificates must be created within the company hosting the BIG-IPs.

ANSWER: A

Explanation:

No SSL certificates are required on the BIG-IP. A Server SSL profile controls the SSL/TLS connection that the BIG-IP establishes *to the pool member or other backend server.* In this role, BIG-IP acts as an SSL client, while the backend server presents its own certificate during the TLS handshake. The Server SSL profile itself does not require BIG-IP to present a local certificate and private key; those are associated with a Client SSL profile, where BIG-IP terminates client-side SSL/TLS connections.

BIG-IP can optionally be configured to validate the certificate presented by the backend server. For example, the Server SSL profile can be configured with trusted certificate authorities and peer-certificate verification settings. In that case, a relevant CA certificate bundle may be installed on BIG-IP so it can verify the server certificate chain. That is a validation configuration choice, however, rather than a certificate intrinsically required simply because a virtual server uses a Server SSL profile. F5 describes Server SSL profiles as enabling SSL encryption for connections from BIG-IP to pool members, and documents the optional server-certificate verification controls in the profile settings. See F5 Server SSL Profile Reference and F5 SSL profile overview.

QUESTION NO: 65

Which two protocol are used for the retrieval of email? (Choose two)

A. ICAP

B. SMTP

C. IMAP

D. POP3

E. SNMP

ANSWER: C D

Explanation:

IMAP and POP3 are the standard application-layer protocols used by email clients to retrieve messages from a mail server. IMAP (Internet Message Access Protocol) is designed for accessing and managing email while it remains stored on the server. It supports mailbox folders, message flags, server-side searching, and synchronization among multiple clients, making it well suited to users who access the same mailbox from several devices. POP3 (Post Office Protocol version 3) enables a mail client to retrieve messages from a server, traditionally by downloading them for local access; clients may be configured to retain copies on the server as needed. Both protocols therefore provide the mail-retrieval function between a user's mail client and a mail server. This distinction is important in application delivery because protocol-aware systems must recognize the relevant client-to-server application traffic and its expected behavior. For protocol specifications, see the IETF's IMAP4rev2 specification (RFC 9051) and POP3 specification (RFC 1939).

QUESTION NO: 66

The F5 Visual Policy Editor (VPE) is unique to the F5 BIG-IP APM module; no other access management tool has this capability. Select the features that the VPE provides. Select two.

A. Determines antivirus patch levels

B. Customizes landing or login page

C. Provides vulnerability scanning

D. Checks operating system patch levels

E. Assigns a lease pool address

ANSWER: A D

Explanation:

The Visual Policy Editor provides a graphical framework for building BIG-IP Access Policy Manager access policies, including endpoint-security checks that evaluate the condition of a client device before access is granted. “Determines antivirus patch levels” is correct because an access policy can use endpoint inspection and antivirus-related checks to collect and evaluate the installed antivirus product, its state, and associated compliance information. This enables an organization to permit, deny, or otherwise steer a session according to endpoint-security requirements.

“Checks operating system patch levels” is also correct because endpoint inspection can gather operating-system and patch-management information from supported client endpoints. A VPE policy can then apply branching logic based on that posture data, allowing access decisions to reflect whether the device satisfies the organization’s required patch level. These checks are implemented as policy items in the visual policy flow, so administrators can combine them with authentication, authorization, and other session logic. F5 documents endpoint inspection as an APM capability for collecting endpoint information and enforcing endpoint security policy, and documents the Visual Policy Editor as the interface for configuring access-policy actions and branches. See the F5 APM endpoint inspection documentation and the F5 Visual Policy Editor documentation.

QUESTION NO: 67

Which FTP mode should be used by a client behind a firewall that has no special configurator?

A. Passive FTP

B. Secure FTP

C. Active FTP

D. Protected FTP

ANSWER: A

Explanation:

Passive FTP is correct because it is designed to work when the client is protected by a firewall or network-address translation device that does not include an FTP-aware configurator. FTP uses a control connection and a separate data connection. In passive mode, the client first establishes the control connection to the server, then requests passive operation. The server supplies a port for data transfer, and the client initiates the data connection outbound to that server port. This means the client-side firewall can apply ordinary outbound connection rules rather than needing to recognize FTP commands and dynamically accommodate a connection initiated back toward the client. This behavior is especially useful for clients located on private networks, where unsolicited inbound data connections cannot be routed or permitted without protocol-specific assistance. The FTP specification describes passive operation through the PASV command and the server-provided data-transfer address and port. F5 also provides an FTP profile specifically to interpret FTP control traffic when protocol-aware handling is required. See RFC 959, FTP passive data transfer and F5 FTP profile documentation.

QUESTION NO: 68

Which of the following is not a method of protection for user-input parameters?

A. Value extraction

B. Attack signatures

C. Length restriction

D. Meta character enforcement

ANSWER: A

Explanation:

Value extraction is the correct answer because it describes obtaining or parsing a value from an HTTP request, rather than applying a security control to restrict what that value may contain. Parameter protection in an F5 Advanced WAF/ASM policy is based on defining and enforcing acceptable parameter characteristics. This includes constraints such as expected data format, permitted characters, and acceptable value size, together with inspection against known attack patterns. These controls allow the policy to evaluate supplied input and block requests that violate the application's defined expectations or match malicious payload patterns.

Extracting a value can be useful to application processing or to identifying data within a request, but extraction alone does not validate, constrain, sanitize, or reject user-controlled content. It therefore is not itself a protection method for user-input parameters. Effective input protection requires positive enforcement of the application's expected input boundaries, which aligns with F5 WAF policy configuration concepts and the general secure-development principle of server-side input validation. See the F5 Technical Documentation portal and OWASP's Input Validation guidance.

QUESTION NO: 69

A GTM System performs a name resolution that is not a Wide-IP. The name is in a domain for which the GTM System is authoritative. Where does the information come from?

A. It comes from BIND database (zone) files on the GTM System.

B. GTM System cannot resolve a host name that is not a Wide-IP.

C. It comes from the database of previously cached name resolutions.

D. It comes from a zone transfer initiated when the request was received.

ANSWER: A

Explanation:

It comes from BIND database (zone) files on the GTM System. BIG-IP DNS, formerly called GTM, can answer DNS requests both for Wide IPs that it manages through its DNS/GTM configuration and for ordinary DNS resource records in zones for which it is authoritative. When a queried name is not configured as a Wide IP but belongs to an authoritative zone, the system uses the zone data maintained by its BIND/named DNS service. That zone data contains standard DNS resource records, such as address, alias, mail-exchanger, and name-server records, and enables the system to provide an authoritative response for conventional host names within the zone.

Administrators manage this type of authoritative DNS zone information through BIG-IP DNS zone management facilities, including ZoneRunner, which works with the BIND service and its zone files. The response is therefore based on the configured authoritative zone database present on the GTM system at the time of the query. This behavior allows a BIG-IP DNS deployment to serve traditional authoritative DNS records alongside its global traffic-management functions for Wide IPs. F5 documentation describes BIG-IP DNS as integrating DNS services with traffic management and documents ZoneRunner management of BIND zones and records. See F5 BIG-IP DNS Concepts and F5 BIG-IP DNS Implementations.

QUESTION NO: 70

Which three must be done so that Generic Host Servers can be monitored using SNMP? (Choose three.)

A. The SNMP monitor must be added to all BIG-IP Systems.

B. The Generic Host Server must be running the big3d agent.

C. The GTM System must be configured for the appropriate MIB.

D. The Generic Host Server must be added to the GTM Configuration.

E. The Generic Host Server must be enabled to answer SNMP queries.

ANSWER: C D E

Explanation:

To monitor a Generic Host Server through SNMP, the GTM configuration needs an object representing that external host, and the host must be reachable and prepared to respond to SNMP polling. Adding the Generic Host Server to the GTM Configuration establishes the target that GTM can associate with its monitoring configuration. The Generic Host Server must also be enabled to answer SNMP queries: an SNMP agent must be running, network access from the GTM system must be permitted, and the configured SNMP version, credentials or community string, and UDP port must be usable by the monitor.

The GTM System must be configured for the appropriate MIB because an SNMP monitor obtains a specific object identifier value and evaluates that value against its configured expectations. The relevant MIB defines the managed objects and object identifiers that identify the health or status data to retrieve. Together, the configured host object, responsive SNMP service, and applicable MIB allow GTM to issue a meaningful SNMP query and use its response for availability decisions. F5 describes monitor configuration and monitor types in its BIG-IP DNS monitor documentation; SNMP protocol and MIB fundamentals are defined in RFC 3418.

QUESTION NO: 71

An administrator connects two devices using an Ethernet cable. The link fails to come up on either device, which setting could prevent the link from being established?

A. Proxy settings

B. IP configuration

C. Link speed

D. DNS resolvers

ANSWER: C

Explanation:

Link speed is correct because Ethernet link establishment begins at the physical interface layer, before IP addressing, name resolution, or application-related settings are used. The two connected interfaces must be able to negotiate, or be configured for, a mutually supported signaling speed. For example, a port forced to a particular speed may not establish a usable physical link with a peer that is configured incompatibly or whose capabilities do not include that speed. This can leave the interfaces reporting that the link is down even when the cable is connected correctly.

On BIG-IP systems, interface status and negotiated link characteristics are fundamental to network connectivity. Administrators should verify the interface media settings, supported speeds, and whether auto-negotiation is enabled consistently at both ends of the connection. When troubleshooting, confirming the physical port state and speed/duplex negotiation is an appropriate early step, since higher-layer configuration is not involved until an Ethernet link is active. The Ethernet standards define the physical-layer and link signaling mechanisms used for these connections; see the IEEE 802.3 Ethernet Working Group. For BIG-IP networking concepts and interface configuration context, consult F5 CloudDocs.