

BIG-IP DNS Specialist

F5 302

Version Demo

Total Demo Questions: 26

Total Premium Questions: 267

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Troubleshooting	53
Topic 2, Design and Implement BIG-IP DNS	38
Topic 3, Configure BIG-IP DNS	169
Topic 4, Mix Questions	7
Total	267

QUESTION NO: 1

Is it a must that the Client performs the “SOA” Record Lookup during the TCP Connection ?

- A. Yes
- B. No

ANSWER: B

Explanation:

No is correct. A DNS client is not required to issue a separate SOA-record lookup as part of establishing or using a TCP connection for a zone transfer. For a full zone transfer, the client opens a TCP connection and sends an AXFR query; the authoritative server's transfer response includes the zone's SOA record at the beginning and again at the end, which provides the serial and boundary information needed for the transfer. Therefore, the protocol does not mandate a distinct preliminary SOA lookup during that TCP connection. In practice, a secondary DNS server may query an SOA record at other times to check a zone's serial number and decide whether a transfer is needed, but that is a refresh/transfer-decision operation rather than a compulsory step in the TCP connection itself. This distinction matters when configuring BIG-IP DNS zone-transfer behavior: TCP is required for AXFR transport, while a separate SOA query is not inherently required by the AXFR transaction. The AXFR protocol requirements are defined in [RFC 5936](#); SOA serial-based incremental-transfer behavior is described in [RFC 1995](#).

QUESTION NO: 2

How does the configuration utility displays the failing Server Object ?

- A. “Offline (Red)”
- B. “online (Red)”
- C. “Offline (green)”

ANSWER: A

Explanation:

“Offline (Red)” is the correct status indication for a failing BIG-IP DNS Server object in the configuration utility. BIG-IP DNS uses the server object's health and iQuery connectivity information to establish whether it can communicate with, synchronize with, and obtain operational data from the associated BIG-IP system. When that server cannot be reached or its relevant health state has failed, the interface represents the object as offline and uses red to make the failed condition immediately visible to the administrator. This visual state is operationally important because BIG-IP DNS relies on server objects for information such as virtual-server availability and wide-IP pool-member selection. An offline server should therefore prompt investigation of iQuery communication, network reachability, device trust/configuration, and the health of the remote BIG-IP system before relying on its status data for DNS traffic-management decisions. F5's BIG-IP DNS documentation describes server objects and their role in gathering status from managed systems; see the [BIG-IP DNS documentation](#) and the [F5 Technical Documentation portal](#).

QUESTION NO: 3

What is the format of the “TXT” record?

- A. It is either unstructured Text Format or Structured Format
- B. It is .exe format
- C. it is .py format

ANSWER: A

Explanation:

It is either unstructured Text Format or Structured Format is correct because a DNS TXT resource record carries one or more character strings and does not impose a single universal syntax on the text's content. TXT data may therefore be free-form, human-readable text, such as descriptive or verification information, or it may follow a defined structured convention understood by a particular application. Common operational uses include domain-ownership verification and email-authentication policies, where systems interpret text according to an agreed specification. For example, SPF policy data is conventionally published in TXT records using a structured sequence of mechanisms and modifiers, while other TXT records can simply contain arbitrary textual information. In BIG-IP DNS/GTM environments, TXT records remain standard DNS resource records; their wire representation and semantics are governed by DNS standards rather than by an executable or source-code file type. RFC 1035 specifies TXT record data as one or more character strings, which provides the flexibility for both unstructured content and application-defined structured content. See [RFC 1035, section 3.3.14](#) and [RFC 7208, SPF records](#).

QUESTION NO: 4

What is the DNSSEC ?

- A. It is a standard protocol
- B. It works as an extension to DNS Protocol
- C. It deletes all DNS Record
- D. It cleans the DNS System

ANSWER: A B**Explanation:**

DNSSEC (Domain Name System Security Extensions) is a set of IETF standards that adds origin authentication and data-integrity protection to DNS data. It does this through public-key cryptography: DNS zone data is signed with DNSKEY and RRSIG records, and a validating resolver verifies those signatures using a chain of trust that begins at a configured trust anchor, normally the DNS root. This lets the resolver establish that the DNS response was published by the authoritative zone and was not altered in transit.

Accordingly, "It is a standard protocol" is correct in the sense that DNSSEC is a standardized DNS security specification, and "It works as an extension to DNS Protocol" is correct because DNSSEC extends DNS with additional resource records and validation behavior while retaining normal DNS resolution mechanisms. In BIG-IP DNS deployments, DNSSEC signing and validation can be used to provide authenticated DNS answers and support secure DNS delegation chains. The IETF's DNSSEC overview is defined in [RFC 4033](#); the DNSSEC protocol record and signing specifications are provided in [RFC 4034](#).

QUESTION NO: 5

What are the functions of the big3d install script ?

- A. It installs the current version of the big3d process on a remote BIGIP System
- B. It copies the trusted device certificate from the local GTM to the "/config/gtm/client.crt" file on remote BIGIP System
- C. None
- D. It installs the Expired version of the big3d process on a remote BIGIP System
- E. It deletes the trusted device certificate from the local GTM to the "/config/gtm/client.crt" file on remote BIGIP System
- F. It copies the trusted server certificate from the remote BIGIP System to the "/config/gtm/server.crt" file on the local GTM

ANSWER: A B F

Explanation:

The big3d install script supports secure and compatible communications between a BIG-IP DNS system and a remote BIG-IP system used as a server. It installs the current version of the `big3d` process on the remote BIG-IP system, ensuring that the remote system runs a version compatible with the BIG-IP DNS device collecting probe and path-performance information. The installation workflow also establishes the mutual certificate trust required for communication between the systems. Specifically, it copies the local GTM system's trusted device certificate to `/config/gtm/client.crt` on the remote BIG-IP system. In the other direction, it retrieves the remote BIG-IP system's trusted server certificate and stores it as `/config/gtm/server.crt` on the local GTM system. These certificate exchanges allow the BIG-IP DNS system and the remote system to authenticate their communications while BIG-IP DNS obtains metrics used for load-balancing decisions. F5 documents the relationship between BIG-IP DNS, BIG-IP servers, and the big3d daemon in its [BIG-IP DNS documentation](#). Additional configuration and operational guidance is available in the [F5 knowledge article portal](#).

QUESTION NO: 6

A Wide IP uses persistence to keep DNS responses consistent for a client network. Two users are located behind the same recursive LDNS server, but they are on different internal IP subnets.

Why can BIG-IP DNS return the same persisted answer to both users?

- A. The users have identical browser cookies
- B. The DNS requests have the same LDNS source address
- C. The users are assigned to the same virtual server VLAN
- D. The Wide IP always returns one address for every query type

ANSWER: B

Explanation:

BIG-IP DNS normally evaluates persistence by using the address of the requesting LDNS server, not the individual end-user address hidden behind that resolver. Since both users send their DNS queries through the same LDNS server, BIG-IP DNS sees the same request source for persistence matching and can return the same persisted resource.

QUESTION NO: 7

What is the GTM Screening mode Deployment type ?

- A. It is named "Authoritative Screening Architecture." It enables the use of all GSLB features.
- B. It allows GTM to receive all DNS queries. It allows GTM to load balance incoming DNS queries to a pool of DNS servers.

ANSWER: A B

Explanation:

"It is named 'Authoritative Screening Architecture.' It enables the use of all GSLB features." is correct because Screening mode places BIG-IP DNS (formerly GTM) in the authoritative DNS resolution path. BIG-IP DNS can inspect requests for the zones it serves and apply Global Server Load Balancing logic, including Wide IP processing, pool selection, load-balancing methods, availability monitoring, persistence, topology decisions, and disaster-recovery behavior. This architecture therefore allows the full DNS/GSLB decision process to be applied when a matching request arrives.

"It allows GTM to receive all DNS queries. It allows GTM to load balance incoming DNS queries to a pool of DNS servers." is also correct. In an authoritative screening deployment, BIG-IP DNS is the DNS-facing device that receives queries. It can answer requests for GSLB-enabled names itself and can direct or distribute DNS traffic for names that require conventional authoritative DNS handling to the appropriate DNS-server resources. This lets the organization retain DNS-server infrastructure while using BIG-IP DNS as the authoritative control point for intelligent global traffic decisions. F5's BIG-IP

DNS documentation describes BIG-IP DNS as the system that responds to DNS requests and applies GSLB configuration and load-balancing decisions; see [F5 BIG-IP DNS Concepts](#) and [F5 BIG-IP DNS Training](#).

QUESTION NO: 8

What are the pre-requests for the BIGIP Systems to share the Data using iQuery?

- A. They should be members of the same synchronization group
- B. They should exchange their SSL digital certificates with each other
- C. They should exchange their neighbors with each other
- D. They should copy their neighbors

ANSWER: A B

Explanation:

For BIG-IP DNS systems to exchange operational data through iQuery, the systems must be configured as members of the same synchronization group and must exchange SSL digital certificates with each other. A synchronization group defines the set of BIG-IP DNS devices that participate together in configuration synchronization and in the associated DNS environment. Membership provides the common topology and configuration relationship needed for the devices to operate as a coordinated group.

iQuery is a secure, proprietary communication channel used by BIG-IP DNS to obtain information such as virtual-server availability, server status, and other metrics needed for intelligent DNS decisions. The peer devices must establish mutual trust before this information can be exchanged. Exchanging SSL certificates provides that trust relationship: each BIG-IP system can authenticate the identity of its iQuery peer and protect the communications channel with TLS. Once the synchronization-group relationship and certificate-based trust are in place, the BIG-IP DNS systems can communicate securely using iQuery and share the state information required for DNS health monitoring and load-balancing decisions.

See the F5 documentation on [communication between BIG-IP DNS systems](#) and the [BIG-IP DNS synchronization process](#).

QUESTION NO: 9

What is the importance of the SOA DNS Record ?

- A. It starts each Zone file
- B. It identifies the primary Authoritative name server of the Zone
- C. It identifies the Mail Server
- D. It identifies the Mx Server

ANSWER: A B

Explanation:

An SOA (Start of Authority) record is mandatory at the apex of every DNS zone and establishes the administrative authority and core operational settings for that zone. Traditionally, it is the first record in a zone file, so “It starts each Zone file” accurately reflects its role in defining the zone. The record’s MNAME field identifies the primary authoritative name server, meaning the server that is the original source for the zone’s data and from which secondary authoritative servers can obtain zone transfers. Therefore, “It identifies the primary Authoritative name server of the Zone” is also correct. In addition to the primary server, an SOA record contains the responsible party’s mailbox and timing values used in zone synchronization and caching, including the serial number, refresh, retry, expire, and negative-cache TTL values. DNS secondaries use these values, particularly the serial number, to determine whether their copy of the zone must be updated. These functions make the SOA record foundational to authoritative DNS zone administration and replication. See [RFC 1035, section 3.3.13](#) and [RFC 1912, section 2.6](#).

QUESTION NO: 10

By Which Organization is the Root Domain Managed?

- A. By the "IANA"
- B. By "IEEE"
- C. By Cisco
- D. By AWS

ANSWER: A

Explanation:

By the "IANA" is correct. The Internet Assigned Numbers Authority (IANA) performs the Internet's root zone management function. The DNS root is the top of the hierarchical namespace and contains delegations for top-level domains such as generic domains and country-code domains. As part of root zone management, IANA processes requests associated with changes to the root zone, including top-level-domain delegation and technical updates, according to established policies and with the required authorization. This function is fundamental to maintaining one globally coordinated DNS namespace, so that DNS resolvers worldwide can begin resolution from a consistent root.

In the current operational model, the IANA functions are carried out by Public Technical Identifiers (PTI), an ICANN affiliate, under agreements and community-developed policies. For exam purposes, IANA remains the recognized organization responsible for coordinating DNS root-zone management. F5 BIG-IP DNS relies on the globally distributed DNS hierarchy and root-server infrastructure for DNS resolution processes, making the IANA role an important foundational DNS concept. See the [IANA Root Zone Management](#) page and ICANN's [IANA Functions](#) overview for details.

QUESTION NO: 11

What is the command that should be issued after installing the SSL Certificate and the private key ?

- A. tmsh save /sys config
- B. ssh save /sys config
- C. tmsh del /sys config

ANSWER: A

Explanation:

`tmsh save /sys config` is the appropriate command because it writes the active BIG-IP configuration to the persistent configuration files. Certificate and private-key objects are configuration items on BIG-IP; after they have been imported or created, saving the configuration ensures that the current running configuration is retained across a system restart or failover event. In operational practice, an administrator should validate that the certificate and key are correctly installed and associated with the intended SSL profile or other consuming object, then save the configuration from tmsh. The command uses the `tmsh save sys config` syntax, where `sys config` identifies the system configuration to persist. This is particularly important when changes were made interactively, because unsaved running-configuration changes are not guaranteed to remain after reboot. F5 documents the `save sys config` command in its tmsh reference and describes certificate and key management as part of BIG-IP system configuration. See the [F5 tmsh sys config reference](#) and [F5 SSL certificate management documentation](#).

QUESTION NO: 12

Using the TMSH

How can we check the Self Ips of the Synchronization group GSLB Objects ?

- A. `tmsh list /gtm server addresses`
- B. `tmsh display /gtm server addresses`

ANSWER: A

Explanation:

The command `tmsh list /gtm server addresses` is the appropriate TMSH query for viewing the addresses configured on BIG-IP DNS (GTM) server objects. In a GSLB synchronization group, each GTM server object stores one or more addresses used to identify and communicate with that server. These configured addresses are the relevant Self IP values to verify when checking the peer definitions associated with the synchronized GSLB configuration.

Running the command without naming an individual server returns the `addresses` property for all GTM server objects, making it useful for reviewing the entire synchronization group and confirming that each member is represented with the expected reachable address. The leading slash is valid TMSH object-path notation and explicitly identifies the GTM module. This query reads the configured object data; it does not modify the GSLB configuration. F5 documents the GTM server object and its `addresses` property in the [TMSH GTM server reference](#). Additional BIG-IP DNS configuration documentation is available in the [BIG-IP DNS configuration guide](#).

QUESTION NO: 13

True or false :

The recursion is disabled by default in F5 GTM System ?

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. BIG-IP DNS, formerly called GTM, is normally deployed as an authoritative DNS and global traffic-management platform. Its DNS listeners answer queries for zones and Wide IP records for which the system is authoritative; they do not act as open recursive resolvers by default. Recursive processing must be explicitly configured when a deployment requires the BIG-IP system to resolve names through upstream DNS servers. Keeping recursion off by default is an important secure DNS design principle because it prevents the appliance from being unintentionally exposed as a public recursive resolver. In particular, unrestricted recursive service can be abused by unauthenticated Internet clients and can increase exposure to reflected or amplified DNS traffic. Where recursive DNS behavior is genuinely required, administrators should configure the relevant resolver or forwarding behavior deliberately and restrict access with appropriate listener, network, and security controls. This default behavior supports the expected GTM/BIG-IP DNS role of responding authoritatively and making intelligent DNS distribution decisions rather than providing general-purpose recursive resolution. F5 documents DNS resolver configuration separately from normal BIG-IP DNS listener and authoritative-service configuration: [F5 BIG-IP DNS: Configuring DNS resolvers](#). F5's BIG-IP DNS documentation portal also provides the product configuration references: [BIG-IP DNS documentation](#).

QUESTION NO: 14

What is the Zone Apex ?

- A. It is the basic domain name of the DNS zone.
- B. It does not include labels for names beneath the zone, such as subdomains or hosts.
- C. It is the node of the DNS namespace at the top of the zone.

ANSWER: A B C

Explanation:

A zone apex is the domain name at the top of a DNS zone—the node from which the zone's authoritative data is defined. For example, if a zone is named `example.com.`, then `example.com.` is its apex; names such as `www.example.com.` are nodes below that apex. Thus, describing it as the basic domain name of the zone and as excluding subordinate labels accurately expresses the relationship between the apex and names beneath it. The zone apex is also specifically a node in the DNS namespace, not a web-specific construct: a DNS zone can serve many DNS uses and need not correspond to a website at all. DNS standards describe a zone as beginning at the top node where its authoritative data is maintained, and current DNS terminology defines the zone apex as that top node. At the apex, the zone commonly contains essential records such as SOA and NS records, which establish the zone's authority and delegation behavior. See the DNS concepts in [RFC 1034](#) and the formal terminology in [RFC 9499](#).

QUESTION NO: 15

At which Location will the DNS Records be stored inside the PC?

- A. There are 2 locations:
- B. PC Cache
- C. Host file
- D. ROM

ANSWER: B C

Explanation:

The correct local locations are the **PC Cache** and the **Host file**. A computer's DNS client maintains a local resolver cache containing DNS resource records that it has recently resolved. When an application requests a name, the client can use an unexpired cached record immediately, avoiding another query to a configured DNS server and reducing lookup latency. Cached entries are temporary: their lifetime is controlled by the record's time-to-live value, after which the DNS client must obtain current information again.

The Host file is also a local name-resolution source. It contains manually maintained hostname-to-address mappings and is consulted as part of local name resolution. This permits an administrator to define a specific mapping locally, such as for testing, troubleshooting, or temporarily directing a name to a particular address. Because the mapping is local to that machine, it applies only on the computer containing the file. Microsoft documents DNS client behavior and its local cache in its [DNS client resolution documentation](#); the Hosts file format and purpose originate from [RFC 952](#).

QUESTION NO: 16

What is the SNMP Manger?

- A. It is the management system from which the SNMP agent will accept requests.
- B. It is the part of the SNMP System that runs on the F5 BIGIP System

ANSWER: A

Explanation:

It is the management system from which the SNMP agent will accept requests. In the SNMP architecture, the manager is the network-management application or system that initiates management operations. It sends requests to an SNMP agent, such as requests to retrieve management information from the agent's Management Information Base, and it can issue authorized configuration requests where the SNMP version and access controls permit them. The manager also receives asynchronous notifications, including traps and informs, sent by configured agents. On a BIG-IP system, SNMP configuration defines the agent behavior, communities or SNMPv3 security settings, permitted management hosts, and notification destinations; these settings govern which management systems can query the device and receive its notifications. This manager-agent relationship is a core SNMP architectural concept: the agent runs on the managed device, while the

manager is external software used to monitor and administer that device. F5's SNMP tmsh documentation describes the BIG-IP SNMP agent configuration, including allowed clients and trap destinations, while the SNMP architecture RFC defines the manager's role in originating requests and processing notifications. See [F5 BIG-IP sys snmp reference](#) and [RFC 3411: SNMP Management Framework](#).

QUESTION NO: 17

A DNSSEC zone is configured to use NSEC3 for authenticated denial of existence.

Which two statements describe NSEC3 behavior? Select two.

- A. It uses hashed owner names in denial-of-existence records.
- B. It makes straightforward zone walking more difficult than NSEC.
- C. It encrypts DNS answers between BIG-IP DNS and the recursive resolver.
- D. It removes the requirement to generate RRSIG records.
- E. It replaces the DS record at a secure delegation.

ANSWER: A B

Explanation:

NSEC3 provides authenticated denial of existence by publishing hashed owner-name information rather than the clear-text owner names used by NSEC. This makes straightforward enumeration of zone names more difficult. NSEC3 does not encrypt DNS responses, and it does not replace the RRSIG records that provide DNSSEC signatures over DNS data.

QUESTION NO: 18

What are the main drawbacks of the Synchronization / iQuery Connection issues ?

- A. There are differences in the configuration among the Synchronization Group members
- B. The Log messages will report different issues related to Synchronization and the iQuery Connections
- C. The Server Configuration objects and components will be marked as unavailable or unknown
- D. The Server Configuration objects and components will be marked as red or blue

ANSWER: A B C D

Explanation:

There are differences in the configuration among the Synchronization Group members because synchronization connectivity is what distributes and maintains the shared BIG-IP DNS configuration across the group. When synchronization cannot complete reliably, members can retain different versions of objects and settings, reducing configuration consistency. The Log messages will report different issues related to Synchronization and the iQuery Connections because BIG-IP DNS records failures to establish, maintain, or use these control-plane communications; these messages are essential indicators of the affected peer, server, or connection state. The Server Configuration objects and components will be marked as unavailable or unknown because BIG-IP DNS uses iQuery to obtain status and metrics from BIG-IP systems. Without a usable iQuery relationship, BIG-IP DNS cannot reliably determine the health and availability information it normally receives from the remote system. The Server Configuration objects and components will be marked as red or blue reflects those unavailable or unknown operational states in the BIG-IP DNS interface. These effects can directly impair topology and load-balancing decisions because DNS responses depend on current server availability and synchronized configuration. F5 describes iQuery as the communication mechanism used by BIG-IP DNS to exchange status information with BIG-IP systems, and its DNS documentation covers synchronization-group operation and monitoring. See [F5 BIG-IP DNS Concepts](#) and [F5 BIG-IP DNS documentation](#).

QUESTION NO: 19

What is the iterative Query ?

- A. This type of Query allows the DNS server to direct the client system to other DNS Server that will help resolving the Query
- B. The DNS Server will redirect the Requester to other DNS Server that may help in resolving the Query
- C. The DNS Receiver will continue to query other servers until getting the resolution or even an error message indicating that the FQDN Doesn't exist

ANSWER: A B

Explanation:

An iterative DNS query is one in which the queried DNS server returns the best information it currently has rather than taking responsibility for completing the entire lookup. When it is not authoritative for the requested name and does not have a cached answer, it commonly supplies a referral containing NS records for a DNS server that is closer to, or authoritative for, the requested domain. The requester then uses that referral to send a subsequent query to the identified server. Thus, "This type of Query allows the DNS server to direct the client system to other DNS Server that will help resolving the Query" and "The DNS Server will redirect the Requester to other DNS Server that may help in resolving the Query" both accurately express the essential referral-based behavior of iterative resolution. Iterative lookups are fundamental to DNS delegation: a resolver can begin at a root server, receive a referral to a top-level-domain server, then receive another referral for the authoritative domain, and finally obtain the requested record. RFC 1034 describes this distinction and the use of referrals during DNS resolution. See [RFC 1034, section 4.3.2](#) and F5's [BIG-IP DNS documentation](#).

QUESTION NO: 20

What are the Contents of the MIB ASCII File ?

- A. Consists of
- B. list of the Data Objects
- C. Data Object Type
- D. Data Object Current Validity
- E. Brief Description about the Purpose of each Object

ANSWER: B C D E

Explanation:

An SNMP MIB ASCII file is a human-readable module that defines the managed objects exposed through SNMP. A list of the data objects is fundamental because the MIB identifies each manageable item by name and object identifier. Data Object Type is also required: the object definition's `SYNTAX` clause specifies the object's ASN.1 data type and permitted representation, allowing an SNMP manager to correctly interpret values returned by an agent.

Data Object Current Validity is represented by the object's `STATUS` clause. This records the definition's lifecycle state, such as `current`, `deprecated`, or `obsolete`, so management software and administrators can determine whether a definition remains supported for use. Brief Description about the Purpose of each Object is supplied through the `DESCRIPTION` clause, which documents the object's intended meaning and use. Together, the object list and these definition clauses make the ASCII MIB usable by both people and SNMP management tools.

The SMIv2 rules defining `SYNTAX`, `STATUS`, and `DESCRIPTION` are specified in [RFC 2578](#). F5's SNMP implementation and MIB usage are documented in the [BIG-IP System Monitoring: Setting Up SNMP](#) guide.

QUESTION NO: 21

At which part of the SOA Does the Zone Transfer depend ?

- A. It depends mainly on The Serial number
- B. It depends mainly on The CARD number

ANSWER: A

Explanation:

It depends mainly on The Serial number. The Start of Authority (SOA) resource record includes a serial field that represents the version of a DNS zone. Secondary authoritative DNS servers query the primary server's SOA record at the configured refresh interval and compare the returned serial value with the serial value of their local zone copy. A higher serial on the primary indicates that the zone has been updated, causing the secondary to request a zone transfer—typically an incremental transfer (IXFR), when supported, or otherwise a full transfer (AXFR). Administrators must increment the SOA serial whenever they modify zone data so that secondary servers can detect the change and obtain the current records. The serial is therefore the SOA component that determines whether a zone transfer is needed; it is not itself the mechanism that transports the zone data. RFC 1035 defines the SOA serial as the unsigned 32-bit version number used by secondary servers to determine whether their copy is current. F5 BIG-IP DNS deployments that host or interact with authoritative zones follow this standard DNS synchronization behavior. See [RFC 1035, SOA resource record](#) and [RFC 1996, DNS NOTIFY](#).

QUESTION NO: 22

What are the topology records Builder components ?

- A. Destination
- B. Filtration character
- C. Request Source

ANSWER: A C

Explanation:

The topology-record builder is based on matching the origin of a DNS request to a defined destination. **Destination** is therefore a required builder component: it identifies the target topology object or location that BIG-IP DNS should prefer when the source criteria match. The companion component, **Request Source**, identifies where the DNS request is considered to originate, such as an LDNS address, subnet, region, country, ISP, or other configured topology entity. BIG-IP DNS evaluates topology records to select the most appropriate destination for the request source, allowing DNS responses to reflect network geography, provider, or other routing policies. A topology record consequently expresses a relationship between a request origin and a destination, rather than merely defining a generic traffic-distribution setting. F5 documents topology load balancing as using topology records to direct DNS requests according to the requester's source and the configured destination criteria. See the [F5 BIG-IP DNS topology documentation](#) and the [F5 BIG-IP DNS product documentation](#) for topology-record configuration concepts.

QUESTION NO: 23

What is the Location of the Core files inside the system if the system configure to save the core files ?

- A. "/var/core " Directory
- B. "/variable/core " Directory
- C. "/txt/core " Directory

ANSWER: A

Explanation:

“/var/core” Directory is the correct location. On BIG-IP systems running TMOS, saved core files are placed in /var/core. A core file is a memory image generated when a process terminates unexpectedly; it gives F5 Support and administrators the information needed to inspect the failed process, identify its state at the time of failure, and troubleshoot the underlying issue. Core-file generation and retention can be controlled by system configuration and available disk space, but when BIG-IP is configured to save these diagnostic files, /var/core is the standard on-box directory to check. Administrators commonly include relevant files from this location when collecting diagnostic information for a support case, while taking appropriate care because core files can be large and may contain sensitive information from process memory. F5’s incident-management documentation identifies core files as a key diagnostic artifact and describes their handling on BIG-IP systems. See the [F5 BIG-IP Incident Management documentation](#) and F5’s [guidance for collecting BIG-IP diagnostic information](#).

QUESTION NO: 24

Regarding to the Sync Settings

What is the Condition that allows the Synchronization group members to communicate and Synchronize properly ?

- A. The all members should be configured to use NTP
- B. Some members should be configured to use NTP
- C. The all members should be configured to use zone

ANSWER: A

Explanation:

The all members should be configured to use NTP is correct because BIG-IP DNS synchronization-group communication depends on members maintaining consistent system time. Synchronization groups use iQuery to exchange configuration and state information among BIG-IP DNS devices. If clocks differ substantially, time-sensitive operations can fail or behave unpredictably, and diagnostics may report iQuery or synchronization problems. Configuring every device in the group to synchronize with Network Time Protocol provides a common, accurate time base and prevents clock drift from disrupting communication.

In practice, each BIG-IP system should use reachable, reliable NTP servers and should show synchronized status before it is added to, or expected to participate reliably in, the synchronization group. This is a group-wide prerequisite: a synchronized clock on every member supports dependable trust-related and control-plane exchanges throughout the group. F5 documents NTP configuration and verification for BIG-IP systems in its [NTP configuration guidance](#); the BIG-IP DNS documentation also describes synchronization groups and iQuery-based inter-device synchronization in the [F5 DNS documentation](#).

QUESTION NO: 25

Regarding to the iQuery Statistics :

What will the following identify

The “Bytes In”?

- A. The amount of data in bytes received by the BIG-IP DNS
- B. over the iQuery connection from the specified server.
- C. The amount of data Backed up

ANSWER: A

Explanation:

The amount of data in bytes received by the BIG-IP DNS is correct because the *Bytes In* counter in iQuery statistics measures inbound traffic received through an iQuery session. iQuery is the control-plane communication mechanism that

BIG-IP DNS uses to exchange status, topology, and related synchronization information with other BIG-IP systems. Consequently, this statistic represents the volume of iQuery data arriving at the local BIG-IP DNS from its peer, expressed in bytes. Administrators can use the value to understand whether an iQuery connection is actively carrying inbound information and to establish a baseline for expected synchronization traffic. When investigating a peer relationship, this counter is meaningful alongside other iQuery connection and message statistics because it directly reflects data received by the local device, rather than application traffic handled for DNS clients. F5 describes iQuery as the protocol used for communication between BIG-IP DNS and other BIG-IP systems in its [BIG-IP DNS iQuery documentation](#). Additional BIG-IP DNS product documentation is available through the [F5 Technical Documentation portal](#).

QUESTION NO: 26

How can we prevent the Entire GTM Sync group from rolling back to the contents of the UCS Archive

That will be Installed on a Sync group member ?

- A. Temporary Disconnect the TMM Switch ports of the new GTM Device from the Production Network
- B. Leave the management interface connected
- C. The third step is to Restore the UCS File
- D. None of the above

ANSWER: A B C

Explanation:

To safely restore a UCS archive to a BIG-IP DNS (GTM) device that belongs to a Sync Group, temporarily disconnect the new device's TMM switch ports from the production network before loading the archive. This isolates the traffic-side interfaces through which the device could participate in the production GTM environment while its saved configuration is being applied. Keeping the management interface connected preserves administrative access, allowing the administrator to monitor the restore process and perform the required follow-up configuration without reconnecting the production-facing interfaces prematurely.

After that isolation is in place, restoring the UCS file is the required next action: it applies the archived system configuration to the replacement or new group member while the device remains protected from influencing the active Sync Group. This sequencing is important because a UCS archive can contain GTM configuration data that must be reviewed and reconciled before the device is returned to normal group participation. F5 documents UCS loading through the `tmsh load sys ucs` command and its associated restore behavior in the [tmsh sys ucs reference](#). Additional BIG-IP DNS configuration and synchronization guidance is available in the [BIG-IP DNS Implementations documentation](#).