

***F5 302 – Study Guide - GTM Specialist
Version 1.0***

Eric Mitchell
Channel SE, East US and Federal
F5 Networks
9/25/2014
e.mitchell@f5.com

This page intentionally left blank for duplex printing.



Overview

Welcome to the 302 - GTM Specialist compiled Study Guide. The purpose of this guide is to help you prepare for the F5 302 - GTM Specialist exam. The contents of this document are based on the 302 - GTM Specialist Blueprint Guide.

This study guide provides students with *some* of the basic foundational knowledge required to pass the exam.

This study guide is a collection of information and therefore not a completely original work. The majority of the information is compiled from F5 sources that are located on the Internet. All of the information locations are referenced at the top of each topic instead of in an Appendix of this document. This was done to help the reader access the referenced information easier without having to search through a formal appendix. This guide also references the same books as the exam Resource Guide for each topic when applicable for consistency. Those books are a great source of information on DNS and Global Traffic Manger (GTM).

F5 Networks provides the 302 - GTM Specialist Resource Guide as a study guide. The Resource Guide is a list of reading material that will help any student build a broad base of general knowledge that can assist in not only their exam success but also in becoming a well-rounded systems engineer. The Resource Guide will be available to the candidate once they are qualified for the GTM Specialist exam.

Taking certified F5 GTM training, such as Configuring BIG-IP GTM v11: Global Traffic Manager, will surely help with the topics of this exam but does not teach directly to the exam content. Hands on administrative experience with the BIG-IP platform licensed with GTM will reinforce many of the topics contained in the 302 - GTM Specialist exam.

The F5 Certified BIG-IP Administrator (F5-CA), which is made up of the 101 - App Delivery Fundamentals and 201 - TMOS Administration exams, stand as a pre-requisite to this exam.

This guide was prepared by an F5 employee but is not an official F5 document and is not supported by F5 Networks.

Reading = Knowledge = Power

Contents

Overview	2
Contents	3
Printed References.....	7
Introduction	8
Objective - 1.01 - Identify resource record types and their purpose including DNSSEC record types	9
Objective - 1.02 - Identify the different zone types and their purpose	16
Objective - 1.03 - Explain the purpose of tools and when to use them	17
Objective - 1.04 - Explain the dataflow of the DNS query process [iterative, recursive, lame delegation, host file, and resolvers]	18
Objective - 1.05 - Distinguish IPv4 versus IPv6 query including differentiating IPv4/6 transport versus IPv4/6 query type and extrapolating when different query types will be used on different transports.....	21
Objective - 1.06 - Given a DNS hierarchical diagram determine what source IP the GTM will receive the query from.....	23
Objective - 1.07 - Identify DNS security concepts and their purpose [DDOS, DNSSEC, AnyCast, DNSFirewall, site validation, iRules, and impacts of floating self-IP versus non-floating self-IP listener].....	24
Objective - 1.08 - Describe data center, server/virtual server, and object monitoring including explanation of resulting object statuses [prober pools, BIG-IP and generic server objects, monitors, etc.].....	27
Objective - 1.10 - Identify applicable iRules events including application to Wide-IP versus Listener	36
Objective - 1.11 - Identify the purpose of GTM tools and when to use them [checkcert, iqdump, etc.].....	40
Objective - 1.12 - Explain how zone transfers work [multi master, master/slave, DNSEXPRESS, incremental/full, updates (notify/expire)]	42

Objective - 1.13 - Given a scenario determine the impact of a custom DNS profile for various types of queries, determine what response will be given and where it will come from.....	45
Objective - 1.14 - Given a scenario with a specific query source IP address and various pool and Wide-IP loading balancing methods and topology rules/regions determine the response that will be given	49
Objective - 1.15 - Explain sync group/iQuery purpose, configuration and basic requirements	50
Objective - 1.16 - Explain the networking requirements of placing devices within a GTM data center object.....	52
Section 2 - Deployment.....	54
Objective - 2.01 Explain when to configure translation addresses for local data center connectivity.....	54
Objective - 2.02 Explain how to configure GTM sync groups and iQuery	55
Objective - 2.03 Given a set of requirements select the appropriate load balancing methods [ex. Wide-IP level, pool level, different types and combinations]	57
Objective - 2.04 Given a scenario select the appropriate deployment type: screening mode, DNS delegation, caching resolver, and DNS 6 to 4	59
Objective - 2.05 Explain how to configure GTM to return non-Wide-IP supported records [ex. MX, SRV, TXT records, etc.].....	63
Objective - 2.06 Given a scenario of specific virtual server status, pool and Wide-IP load balancing settings determine the answer returned [Single pool versus multiple pools, effect of secondary and fall-back mechanisms in the first pool, effect of topology and topology records at the Wide-IP level versus pool level, and iRule effects]	65
Objective - 2.07 Given a set of topology requirements configure a deployment using user defined topology prefixes	66
Objective - 2.08 Given a scenario configure a deployment using auto-discovery [behavior of delete versus no-delete with auto-discovery, compatibility with translation, and route domains].....	68
Objective - 2.09 Explain the necessary steps and tools to add a new LTM to a sync group	70

Objective - 2.10 Explain the necessary steps and tools to add a new GTM to an existing sync group.....	73
Objective - 2.11 Explain how to troubleshoot and verify sync group mesh	74
Objective - 2.12 Explain the use of device certificates in iQuery [SSL components, expiration, 3rd party certs]	80
Objective - 2.13 Explain how to verify listener responses	81
Objective - 2.14 Explain how to verify that DNSSEC is working	83
Objective - 2.15 Given a scenario explain how to validate system health for proper operation.....	85
Section 3 – Operations and Troubleshooting	87
Objective - 3.01 Given a scenario determine the impact of software updates in a group on monitoring and configuration state.....	87
Objective - 3.02 Given a scenario determine what is the effect of changing the features enabled in a DNS profile	89
Objective - 3.03 Explain how to renew device certificates and update them in the sync group.....	93
Objective - 3.04 Explain the impact of restoring a UCS on a GTM	95
Objective - 3.05 Explain the importance of running compatible versions of big3d on the LTM and GTM.....	97
Objective - 3.06 Explain how to properly add/remove device from iQuery mesh	100
Objective - 3.07 Explain the effect of adding a resource record without using ZoneRunner.....	101
Objective - 3.08 Explain the effects and implications of securing/hardening with respect to normal operation, iQuery and resolution	103
Objective - 3.09 Identify GTM specific command line tools and TMSH GTM specific commands.....	105
Objective - 3.10 Given a scenario determine what information needs to be provided when making a support call	109
Conclusion.....	112



Printed References

These referenced books are and important and should be considered basic reading material for this exam.

(Ref:1) Kozierok, Charles M. 2005. The TCP/IP Guide. No Starch Press, Inc. San Francisco, CA. 94103. ISBN 1-59327-047-X pp 947 -1080

(Ref:2) Liu, Cricket and Albitz, Paul. 2006. DNS and BIND, Fifth Edition. O'Reilly Media, Inc. Sebastopol, CA. 95472. ISBN 978-0-596-10057-5

(Ref:3) Configuring GTM v11 Global Traffic Manager. March 2013 v11.3.0. Edition. F5 Networks Training Course. (Configuring GTM: Module X)

Introduction

Global Traffic Manager Introduction

The F5 BIG-IP Global Traffic Manager (GTM) system intelligently resolves names into IP addresses providing intelligent wide area application traffic management and high availability of IP applications/services running across multiple data centers. GTM adds intelligence to DNS. Mastering GTM requires an in-depth knowledge of DNS. Much of this study guide will refer to the basic workings and functional pieces of DNS. Although the Exam Blueprint is not written in a structure that presents topics in an educational order, it does provide all of the necessary building blocks. The Certified GTM training class from F5 will help with many of the scenario based topics on the test.

Objective - 1.01 - Identify resource record types and their purpose including DNSSEC record types

1.01 – Identify resource record types and their purpose

[Link to Online Topic Content](#)

Resource records are the files that contain details about a zone. These resource records, in a hierarchical structure, make up the domain name system (DNS).

Note: Although case is preserved in names and data fields when loaded into the name server, comparisons and lookups in the name server database are not case-sensitive.

The following resource records types are the most common although there are others:

SOA (Start of authority)

The start of authority resource record, SOA, starts every zone file and indicates that a name server is the best source of information for a particular zone. The SOA record indicates that a name server is authoritative for a zone. There must be exactly one SOA record per zone. Unlike other resource records, you create a SOA record only when you create a new master zone file.

A (Address)

The Address record, or A record, lists the IP address for a given host name. The name field is the host's name, and the address is the network interface address. There should be one A record for each IP address of the machine.

AAAA (IPv6 Address)

The IPv6 Address record, or AAAA record, lists the 128-bit IPv6 address for a given host name.

CNAME (Canonical Name)

The Canonical Name resource record, CNAME, specifies an alias or nickname for the official, or canonical, host name. This record must be the only one associated with the alias name. It is usually easier to supply one A record for a given address and use CNAME records to define alias host names for that address.

DNAME (Delegation of Reverse Name)

The Delegation of Reverse Name resource record, DNAME, specifies the reverse lookup of an IPv6 address. These records substitute the suffix of one domain name with another. The

DNAME record instructs GTM (or any DNS server) to build an alias that substitutes a portion of the requested IP address with the data stored in the DNAME record.

HINFO (Host Information)

The Host Information resource record, HINFO, contains information on the hardware and operating system relevant to GTM (or other DNS).

MX (Mail Exchanger)

The Mail Exchange resource record, MX, defines the mail system(s) for a given domain.

NS (Name Server)

The name server resource record, NS, defines the name servers for a given domain, creating a delegation point and a subzone. The first name field specifies the zone that is served by the name server that is specified in the name servers name field. Every zone needs at least one name server.

PTR (Pointer)

A name pointer resource record, PTR, associates a host name with a given IP address. These records are used for reverse name lookups.

SRV (Service)

The Service resource record, SRV, is a pointer that allows an alias for a given service to be redirected to another domain. For example, if the fictional company Site Request had an FTP archive hosted on archive.siterequest.com, the IT department can create an SRV record that allows an alias, ftp.siterequest.com to be redirected to archive.siterequest.com.

TXT (Text)

The Text resource record, TXT, allows you to supply any string of information, such as the location of a server or any other relevant information that you want available.

The following resource record types are associated with DNSSEC:

DNSKEY

These records contain the public key for the zone. They come in two flavors, a Zone Signing Key (ZSK) and a Key Signing Key (KSK). Generally, the KSK signs only certain records within the zone, while the ZSK signs all of the records. You may have as many of each as required for key-rollover protocols or for your needs.

RRSIG – Resource Record Signature

These records hold the signatures for a specific record type. For instance, you will see an RRSIG for NS records, one for DNSKEY records, etc. One RRSIG record will be generated per ZSK, typically, and for certain records one for each KSK as well.

Note: there is one signature per-key per-RRSET, not per RR.

NSEC – Next Secure

This record is used in “negative answers” to prove that a name does not exist. Each name in a zone has an NSEC record added when signed to allow both positive (this name exists) answers and negative answers (this name does not exist) to be cryptographically secure.

NSEC3 – Next Secure (version 3)

This record is used in “negative answers” to prove that a name does not exist. It is similar in function to the NSEC record, but has some advantages in certain situations.

Zones signed with NSEC are “walkable.” This means the entire contents of a zone can be retrieved simply by following the NSEC chain. Also, every name within a zone must be signed and have NSEC records.

NSEC3 uses cryptographic hashes to prevent zone walking while retaining the ability to prove negative answers. It also allows for an opt-out signing method where only certain names within a zone are signed. For very large zones this opt-out is useful.

DS – Delegation Signer

These are records that are submitted to your zone's parent. They are included only in the parent, and correspond to NS records in that they provide a link between your parent and your zone. They are part of the DNSSEC chain of trust from your zone's parent to your zone.

Because many parent zones are not yet signed, DLV may be used to provide others with a trusted relationship to your zone when your parent is not signed or not prepared to accept DS record submissions.

DLV – DNSSEC Look-aside Validation

These records are in most ways identical to DS records. The only difference is the name on the DLV record. A DS record has your zone's name (example.com) while a DLV record has an additional name (example.com.dlv.isc.org.)

1.01 – Identify DNSSEC purpose and GTM implementation

[Link to Online Topic Content](#)

What it is:

Domain Name System Security Extensions (DNSSEC) is an industry-standard protocol that functions as an extension to the Domain Name System (DNS) protocol. GTM uses DNSSEC to guarantee the authenticity of DNS responses and to return Denial of Existence responses thus protecting your network against DNS protocol and DNS server attacks.

[Link to Online Topic Content](#)

How it works:

DNSSEC keys and zones

GTM responds to DNS requests to a specific zone by returning signed name server responses based on the currently available generations of a key. Before you can configure GTM to handle name server responses that are DNSSEC-compliant, you must create DNSSEC keys and zones.

There are two kinds of DNSSEC keys: zone-signing keys and key-signing keys. GTM uses a zone-signing key to sign all of the records in a DNSSEC record set, and a key-signing key to sign only the DNSKEY record of a DNSSEC record set.

F5 Networks recommends that for emergency rollover purposes, when you create a key, you create a duplicate version of the key with a similar name, but do not enable that version. For example, create a key-signing key called ksk1a that is enabled. Then create a duplicate key, but name it ksk1b, and change the state to disabled. When you associate both of these keys with the same zone, you are prepared to easily perform a manual rollover of the key, if necessary.

In order for GTM to use the keys that you create to sign requests, you must assign the keys to a zone. DNSSEC zones are containers that map a domain name to a set of DNSSEC keys that the system uses to sign DNSSEC-compliant name server responses to DNS queries.

When you create a DNSSEC zone, you must assign at least one enabled zone-signing and one enabled key-signing key to the zone before the GTM can sign requests to that zone.

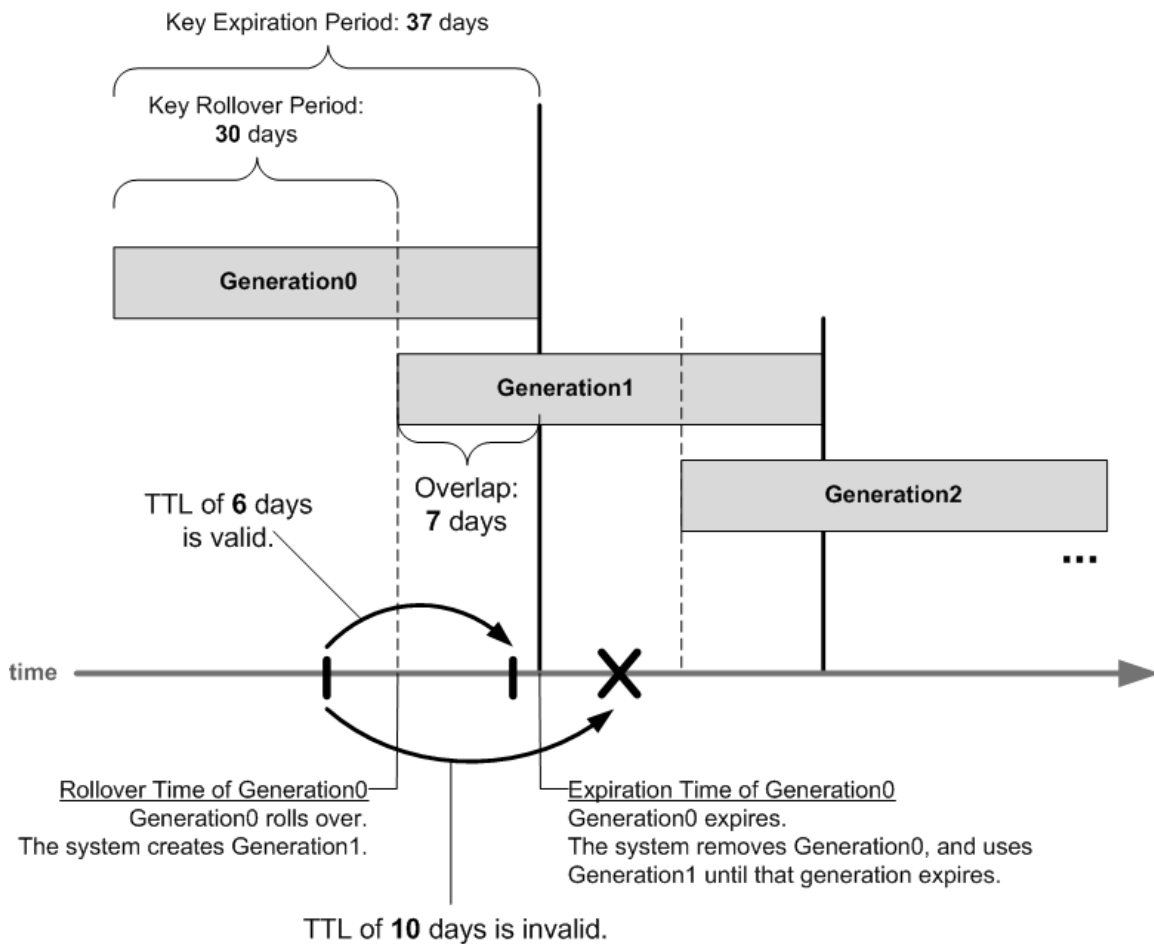
Additionally, after you create a DNSSEC zone, you must submit the DS record for the zone to the administrators of your parent zone, who sign the DS record with their own

key and upload it to their zone. You can find the DS record for your zone in `/config/gtm/dsset-<dnssec.zone.name>`.

Automatic key rollover

To enhance key security, the BIG-IP system has an automatic key rollover feature that uses overlapping generations of a key to ensure that the system can always respond to requests with a signature. The system dynamically creates new generations of each key based on the values of the Rollover Period and Expiration Period settings of the key. The first generation of a key has an ID of 0 (zero). Each time the system dynamically creates a new generation of the key, the ID increments by 1. When a generation of a key expires, the system automatically removes that generation of the key from the configuration.

The following diagram illustrates this, and shows how over time each generation of a key overlaps the previous generation of the key.



Overlapping generations of a key and TTL value

The value that you assign to the TTL (time-to-live) setting for a key specifies how long a client resolver can cache the key. As shown in the figure, the value you assign to the TTL setting of the key must be less than the difference between the values of the Rollover Period and Expiration Period settings of the key; otherwise, a client can make a query and the system can send a valid key that the client cannot recognize.

Important: To ensure that each GTM system is referencing the same time when generating keys, you must synchronize the time setting on each system with the Network Time Protocol (NTP) servers that GTM references.

Providing DS records to the parent domain

Each time a new generation of a key-signing key is created, you must provide the updated DS record to the administrators of the parent zone. For example, in Figure 10.1, the value of the Rollover Period of the key is 30 days, and the value of the Expiration Period of the key is 37 days. In the case of a key-signing key, a new generation of the key is created every 30 days, and you have seven days before the old generation of the key expires to provide the new DS record to the administrators of the parent zone. These administrators sign the new DS record with their own key and upload it their zone.

There are numerous ways to provide the new DS record to the administrators of the parent zone, including secure FTP or use of a secure web site for this purpose. Provide the new DS record to the administrators of the parent zone according to your company policy.


DNSSEC resource records

Your configuration of BIND is independent of the configuration of DNSSEC on GTM. If you want to use BIND for delegation or other tasks, you must add the DNSSEC resource records to your BIND configuration; otherwise, BIND is not aware of these records. If you do this, you can view the DNSSEC resource records in Zone Runner.

[Link to Online Topic Content](#)

GTM Implementation:

You can use GTM to ensure that all responses to DNS-related traffic comply with the DNSSEC security protocol. To configure DNSSEC compliance, you create DNSSEC key-signing and zone-signing keys and a DNSSEC zone. Then you assign at least one enabled key-signing key and one enabled zone-signing key to the zone.



Traffic flow when GTM is the DNSSEC authoritative name server

[Link to Online Topic Content](#)

This link is to the F5 DNSSEC GTM Deployment Guide

Objective - 1.02 - Identify the different zone types and their purpose

1.02 - Identify the different types of zones (Master, Slaves, Hint, Root, Stub)

[Link to Online Topic Content](#)

Zone Types and their Description

Primary (Master) - Zone files for a primary zone contain, at minimum, the start of authority (SOA) and name server (NS) resource records for the zone. Primary zones are authoritative, that is, they respond to DNS queries for the domain or sub-domain. A zone can have only one SOA record, and must have at least one NS record.

Secondary (Slave) - Zone files for a secondary zone are copies of the principal zone files. At an interval specified in the SOA record, secondary zones query the primary zone to check for and obtain updated zone data. A secondary zone responds authoritatively for the zone provided that the zone data is valid.

Stub - Stub zones are similar to secondary zones, except that stub zones contain only the NS records for the zone. Note that stub zones are a specific feature of the BIND implementation of DNS. F5 Networks recommends that you use stub zones only if you have a specific requirement for this functionality.

Forward - The zone file for a forwarding zone contains only information to forward DNS queries to another name server on a per-zone (or per-domain) basis.

Hint - The zone file for a hint zone specifies an initial set of root name servers for the zone. Whenever the local name server starts, it queries a root name server in the hint zone file to obtain the most recent list of root name servers.

Root - The DNS root zone is the top-level DNS zone in the hierarchical namespace of the Domain Name System (DNS) of the Internet. Every name resolution either starts with a query to a root server, or, uses information that was once obtained from a root server. Even if the IP addresses of some root servers change over the years, at least one is needed to retrieve the current list of all name servers. This address file is called named.cache in the BIND name server reference implementation.

Objective - 1.03 - Explain the purpose of tools and when to use them

1.03 - Explain the purpose of tools and when to use them, specifically nslookup, dig, named-checkzone, rndc

[Link to Online Topic Content](#)

DNS Tools

nslookup - nslookup is a tool for querying the Domain Name System (DNS) to obtain domain name or IP address mapping or for any other specific DNS record.

[Link to Online Topic Content](#)

dig – dig is a tool for querying the Domain Name System (DNS) to obtain domain name or IP address mapping or for any other specific DNS record. Dig is part of the BIND domain name server software suite and is available on most linux-based systems.

[Link to Online Topic Content](#)

named-checkzone – named-checkzone checks the syntax and integrity of a zone file. It performs the same checks as named does when loading a zone. This makes named-checkzone useful for checking zone files before configuring them into a name server.

[Link to Online Topic Content](#)

named-compilezone - named-compilezone is similar to named-checkzone, but it always dumps the zone contents to a specified file in a specified format. Additionally, it applies stricter check levels by default, since the dump output will be used as an actual zone file loaded by named. When manually specified otherwise, the check levels must at least be as strict as those specified in the named configuration file.

[Link to Online Topic Content](#)

rndc - BIND includes a utility called rndc which allows command line administration of the named daemon from the localhost or a remote host. rndc controls the operation of a name server. It supersedes the ndc utility that was provided in old BIND releases. If rndc is invoked with no command line options or arguments, it prints a short summary of the supported commands and the available options and their arguments.

rndc reads a configuration file to determine how to contact the name server and decide what algorithm and key it should use.

Objective - 1.04 - Explain the dataflow of the DNS query process [iterative, recursive, lame delegation, host file, and resolvers]

1.04 - Explain the dataflow of the DNS query process

Ref: 1, pp. 1046.

DNS query process

The client system is trying to connect to a remote system via a fully qualified host name; an example would a browser connection to `www.google.com`. The client system will try to resolve the name `www.google.com` to an IP address. The following steps will explain the resolution process.

1. The client system checks its cache to see if it has an address for this name. If it does it will connect to that address. If not the system will proceed to step 2. It could have entries in its cache in two ways:
 - A host file resource record entry
 - Resource records obtained in answered responses from previous DNS queries
2. The client checks its local host file for the name. If there is a match; it resolves the name using this information. If there is no match in the file; the system will proceed to step 3.
3. The client makes a recursive query to its defined local DNS server.
4. The local DNS server receives the request and checks its cache. If it has a matching entry it will return the record info to the client. If there is no match in the local DNS server's cache; the local DNS server will proceed to step 4.
5. The local DNS server generates an iterative request for the name and sends it to a root name server.
6. The root name server does not resolve the name. It returns the name and address of the name server for the ".com" domain.
7. The local DNS server generates an iterative request and sends it to the name server for ".com".
8. The name server for ".com" returns the name and address of the name server for the "google.com" domain.
9. The local DNS server generates an iterative request and sends it to the name server for "google.com".
10. The name server for "google.com" is authoritative for "www.google.com". It returns the IP address for that host to the local DNS server.

11. The local DNS server caches this resolution. (Note: that it will probably also cache some of the other name server resolutions that it received in steps #6, #8 and possible others if there were sub domain requests involved.)
12. The local DNS server returns the name resolution to the client.
13. The client caches the information.
14. Your browser commences an HTTP request to the www.google.com machine's IP address.

[Link to Online Topic Content](#)

Lame Delegation or Lame Response

When performing recursion, the process of looking up a record from the DNS, a name server must generally query several servers, follow up on referrals, and walk down the chain of authority to find the answer.

For each query, the recursing name server expects the other name server to be authoritative for a given zone. For example, the root servers are expected to be authoritative for the root zone. The root servers give out a referral for .com, pointing to a set of servers; any such server is expected to be authoritative for com. The expected authority can be obtained either from a referral for that zone from a parent zone, or from the authority records returned by another authoritative name server for the zone.

If a query is answered in a way that indicates that the responder is not authoritative for the expected zone, the result is called lame. Since the response is almost always in the form of a referral (a delegation response) for either some zone higher up on the tree or for the expected zone itself, the response can be called a lame delegation or lame referral.

1.04 - Explain recursive versus iterative

[Link to Online Topic Content](#)

Recursive vs. Iterative

With a recursive name query, the DNS client requires that the DNS server respond to the client with either the requested resource record or an error message stating that the record or domain name does not exist. The DNS server cannot just refer the DNS client to a different DNS server.

Thus, if a DNS server does not have the requested information when it receives a recursive query, it queries other servers until it gets the information, or until the name query fails.

A DNS client generally makes recursive name queries to a DNS server, or by a DNS server that is configured to pass unresolved name queries to another DNS server, in the case of a DNS server configured to use a forwarder.

An iterative name query is one in which a DNS client allows the DNS server to return the best answer it can give based on its cache or zone data. If the queried DNS server does not have an exact match for the queried name, the best possible information it can return is a referral (that is, a pointer to a DNS server authoritative for a lower level of the domain namespace). The DNS client can then query the DNS server for which it obtained a referral. It continues this process until it locates a DNS server that is authoritative for the queried name, or until an error or time-out condition is met.

[Link to Online Topic Content](#)

GTM Configuration

When a DNS server sets the recursion available (RA) bit in a DNS response, the DNS server is indicating to the client that it will query other name servers for requested domain names if the DNS server's zone files do not contain the answer. By default, DNS recursion is disabled on GTM systems. Under certain circumstances, you may want to enable DNS recursion on the GTM system.

Objective - 1.05 - Distinguish IPv4 versus IPv6 query including differentiating IPv4/6 transport versus IPv4/6 query type and extrapolating when different query types will be used on different transports

1.05 - Explain the difference between IPv6 and IPv4 data transport

[Link to Online Topic Content](#)

The Problem of Name Space Fragmentation - Following the Referral Chain

A resolver that tries to look up a name starts out at the root, and follows referrals until it is referred to a name server that is authoritative for the name. If somewhere down the chain of referrals it is referred to a name server that is only accessible over a transport (IPv4 stack or IPv6 stack) which the resolver cannot use, the resolver is unable to finish the task.

When the Internet moves from IPv4 to a mixture of IPv4 and IPv6; it is only a matter of time until problems with the referral chain start to happen. The complete DNS hierarchy then starts to fragment into a graph where authoritative name servers for certain nodes are only accessible over a certain transport. The concern is that a resolver using only a particular version of IP and querying information about another node using the same version of IP can not do it because somewhere in the chain of servers accessed during the resolution process, one or more of them will only be accessible with the other version of IP.

With all DNS data only available over IPv4 transport everything is simple. IPv4 resolvers can use the intended mechanism of following referrals from the root and down while IPv6 resolvers have to work through a "translator", i.e., they have to use a recursive name server on a so-called "dual stack" host as a "forwarder" since they cannot access the DNS data directly.

1.05 - Explain the difference between IPv6 record and IPv6 data transport

[Link to Online Topic Content](#)

[Link to Online Topic Content](#)

Differences between IPv6 DNS Record and IPv6 DNS Transport.

The IP protocol version used for querying resource records is independent of the protocol version of the resource records; e.g., IPv4 transport can be used to query IPv6 records and vice versa.

Transport has to do with how the requesting (resolving) server sends the request to a DNS server and if that DNS server supports the IP protocol that the requesting server understands (IPv4 or IPv6). If DNS server does not support the IP protocol that the requesting is able to process the request will fail or if it is referred to another name server that does not support that IP protocol the request will fail. To fix the issue the requesting server will have to use a DNS server that has a “dual stack” (supports IPv4 and IPv6) to act as a translator and be a “forwarder” (make the request for the requesting server).

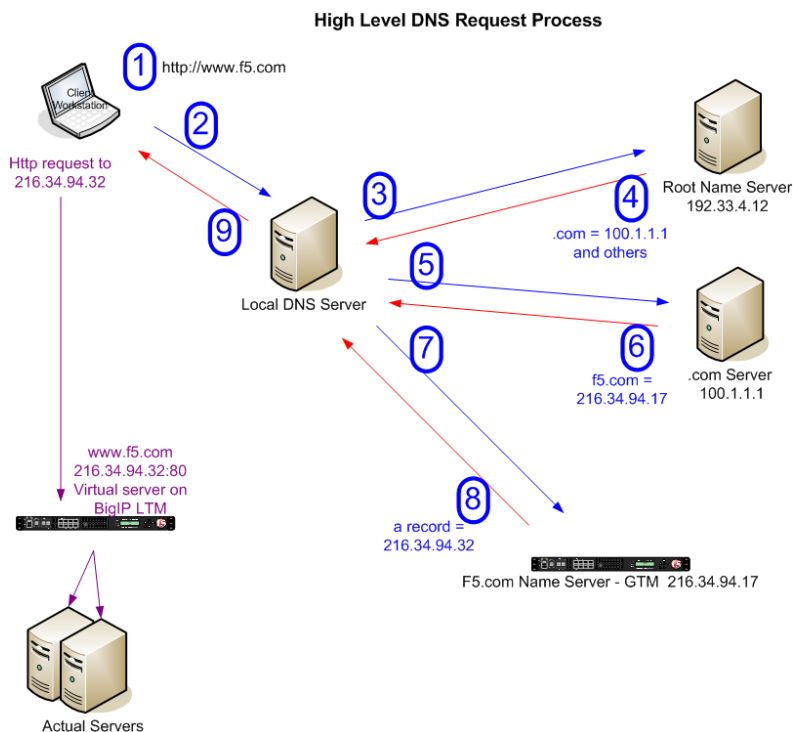
Objective - 1.06 - Given a DNS hierarchical diagram determine what source IP the GTM will receive the query from

1.06 - Given a DNS hierarchical diagram determine what source IP the GTM will receive the query from

Diagram Based Questions

This topic is focused on applying the knowledge you know of the DNS query process. Knowing the process described in section 1.04 will give the candidate the ability to answer the questions on this topic. All F5 exams use diagram-based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

An example of a DNS request environment will be presented and you will need to understand how to overlay the steps described in section 1.04 to the diagram to see which system is actually making the DNS request to the GTM system. The following diagram shows an example of a diagram with the overlaid request process.



Objective - 1.07 - Identify DNS security concepts and their purpose [DDoS, DNSSEC, AnyCast, DNSFirewall, site validation, iRules, and impacts of floating self-IP versus non-floating self-IP listener]

1.07 - Identify DNS security concepts and their purpose [DDoS, DNSSEC, AnyCast, DNSFirewall, site validation, iRules, and impacts of floating self-IP versus non-floating self-IP listener]

[Link to Online Topic Content](#)

Denial of Service (DoS)

DoS and DDoS attacks

Denial of service (DoS) and distributed denial of service (DDoS) attacks attempt to render a machine or network resources unavailable to users. Denial of service attacks require the efforts of one or more people to disrupt the services of a host connected to the Internet. The Advanced Firewall Module allows you to configure packet limits, percentage increase thresholds, and absolute rate limits for a wide variety of packets that attackers leverage as attack vectors, to detect and prevent attacks of this type.

DNS flood (DoS) attacks

Denial of service (DoS) or flood attacks attempt to overwhelm a system by sending thousands of requests that are either malformed or simply attempt to overwhelm a system using a particular DNS query type or protocol extension. The BIG-IP system allows you to track such attacks.

Malformed DNS packets

Malformed DNS packets can be used to consume processing power on the BIG-IP system, ultimately causing slowdowns like a DNS flood. The BIG-IP system drops malformed DNS packets, and allows you to configure how you track such attacks.

Protocol exploits

Attackers can send DNS requests using unusual DNS query types or opcodes. The BIG-IP system can be configured to allow or deny certain DNS query types, and to deny specific DNS opcodes. When you configure the system to deny such protocol exploits, the system tracks these events as attacks.

[Link to Online Topic Content](#)

DNSSEC

DNSSEC is a suite of extensions that add security to the DNS protocol by enabling responses to be validated. The premise behind DNSSEC is that responses to DNS queries need to be trustable. DNSSEC applies the principle of signatures via public/private key encryption as a means to achieve that trust. Essentially DNSSEC is the wrapping of the DNS infrastructure within a trusted, PKI-based superstructure that validates through certificates managed records (zones).

[Link to Online Topic Content](#)

AnyCast

IP Anycast for DNS services on the BIG-IP system to help mitigate distributed denial-of-service attacks (DDoS), reduce DNS latency, improve the scalability of your network, and assist with traffic management. This configuration adds routes to and removes routes from the routing table based on availability. Advertising routes to virtual addresses based on the status of attached listeners is known as Route Health Injection (RHI).

[Link to Online Topic Content](#)

DNSFirewall

A DNS firewall shields DNS infrastructure from attacks such as reflection or amplification DDoS attacks and other undesired DNS queries and responses that reduce DNS performance. In addition, you can mitigate complex DNS security threats by blocking access to malicious IP domains with Response Policy Zones. With GTM, you can install a third-party domain filtering service such as SURBL or Spamhaus and prevent client infection or intercept infected responses to known sources of malware and viruses. F5 DNS firewall services reduce the costs of infection resolution and increase user productivity.

[Link to Online Topic Content](#)

Site Validation

Ability to ensure valid information from known systems through the use of digitally signed answers. This function is a piece of the specifications of DNSSEC.

[Link to Online Topic Content](#)

Impacts of Floating vs Non-Floating Self-IP Listener

A listener is a specialized virtual server that uses port 53 and to which you assign a specific IP address. When a DNS name resolution request is sent to the IP address of a listener, GTM either handles the request locally or forwards the request to the appropriate resource.

When GTM receives traffic, processes it locally, and sends the appropriate Domain Name System (DNS) response back to the querying server, it is operating in Node mode. In this situation, you create a listener that corresponds to an IP address on the system. If GTM operates as a standalone unit, this IP address is the self-IP address of the system. If GTM is part of a redundant system configuration for high availability purposes, this IP address is the floating IP address that belongs to both systems.

Having GTM deployed in a redundant pair configuration in each data center will allow for hitless code upgrades for the GTM listeners residing in that data center. As one GTM is taken off line for maintenance the redundant unit in the pair will continue to respond to the floating IP addresses used as listeners. For any GTM systems that are running stand alone, all listeners on the system will be down during maintenance.

Objective - 1.08 - Describe data center, server/virtual server, and object monitoring including explanation of resulting object statuses [prober pools, BIG-IP and generic server objects, monitors, etc.]

1.08 - Describe data center, server/virtual server, and object monitoring including explanation of resulting object statuses [prober pools, BIG-IP and generic server objects, monitors, etc.]

[Link to Online Topic Content](#)

Data Center Objects

All of the resources on your network are associated with a data center. The GTM consolidates the paths and metrics data collected from the servers, virtual servers, and links in the data center, and uses that data to conduct load balancing and route DNS name resolution requests to the best-performing site based on different factors.

The GTM might send all requests to one site when another site is down. Alternatively, GTM might send a request to the data center that has the fastest response time. A third option might be for GTM to send a request to the data center that is located closest to the client's source address.

Server and Virtual Servers

A server defines a physical system on the network. Servers contain the virtual servers that are the ultimate destinations of DNS name resolution requests. GTM supports three types of servers:

- BIG-IP systems

Any member of the BIG-IP system product line.

- Third-party load balancing systems

A third-party load balancing system is any system, other than a BIG-IP system, that supports and manages virtual servers on the network.

- Third-party host servers

A third-party host server is any server on the network that does not support virtual servers.

A virtual server is a specific IP address and port number that points to a resource on the network. In the case of host servers, this IP address and port number likely point to the resource itself. With load balancing systems, such as the Local Traffic Manager (LTM), virtual servers are often proxies that allow the load balancing server to manage a resource request across a multitude of resources.

[Link to Online Topic Content](#)

Monitors

The GTM health monitors verify connectivity to virtual server objects that reside on BIG-IP devices and non-BIG-IP devices. When a GTM monitor marks a virtual server down, and that virtual server is a member of a Wide-IP pool, the GTM no longer includes that virtual server's address in answers to DNS queries for a Wide-IP.

The GTM system receives the status of virtual servers that reside on non-BIG-IP devices directly from BIG-IP devices using the iQuery protocol. For example, if you have GTM devices and LTM devices in a data center, along with generic host servers, the GTM system will dynamically assign a BIG-IP device to probe the generic host server for status. You can also assign a specific BIG-IP device or group of devices to monitor a server object by using prober pools.

1.08 - Identify the purpose and uses of prober pools

[Link to Online Topic Content](#)


Prober pools

A Prober pool is an ordered collection of one or more BIG-IP systems. GTM can be a member of more than one Prober pool, and a Prober pool can be assigned to an individual server or a data center. When you assign a Prober pool to a data center, by default, the servers in that data center inherit that Prober pool.

The members of a Prober pool perform monitor probes of servers to gather data about the health and performance of the resources on the servers. GTM makes load balancing decisions based on the gathered data. If all of the members of a Prober pool are marked down, or if a server has no Prober pool assigned, GTM reverts to a default intelligent probing algorithm to gather data about the resources on the server.

This figure illustrates how Prober pools work. GTM contains two BIG-IP Local Traffic Manager™ (LTM™) systems that are assigned Prober pools and one LTM system that is not assigned a Prober pool:

Example illustration of how Prober pools work



BIG-IP systems with prober pools

Prober Pool 1 is assigned to a generic host server

BIG-IP LTM3 is the only member of Prober Pool 1, and performs all HTTPS monitor probes of the server.

Prober Pool 2 is assigned to generic load balancers

BIG-IP LTM1 and BIG-IP LTM2 are members of Prober Pool 2. These two systems perform HTTP monitor probes of generic load balancers based on the load balancing method assigned to Prober Pool 2.

The generic load balancers on the left side of the graphic are not assigned a Prober pool

GTM can solicit any BIG-IP system to perform FTP monitor probes of these load balancers, including systems that are Prober pool members.



Objective - 1.09 - Define the GTM load balancing methods and when to use them [dynamic, static]

1.09 - Define the GTM load balancing methods and when to use them [dynamic, static]

[Link to Online Topic Content](#)

Global Server Load Balancing

GTM provides tiered global server load balancing (GSLB). GTM distributes DNS name resolution requests, first to the best available pool in a Wide-IP, and then to the best available virtual server within that pool. GTM selects the best available resource using either a static or a dynamic load balancing method. Using a static load balancing method, GTM selects a resource based on a pre-defined pattern. Using a dynamic load balancing method, GTM selects a resource based on current performance metrics collected by the big3d agents running in each data center.

Static load balancing methods

This table describes the static load balancing methods available in GTM.

NAME	DESCRIPTION	RECOMMENDED USE
Drop Packet	GTM drops the DNS request.	Use Drop Packet for the Alternate load balancing method when you want to ensure that GTM does not offer in a response a virtual server that is potentially unavailable.
Fallback IP	GTM distributes DNS name resolution requests to a virtual server that you specify. This virtual server is not monitored for availability.	Use Fallback IP for the fallback load balancing method when you want GTM to return a disaster recovery site when the preferred and alternate load balancing methods do not return an available virtual server.

Global Availability	GTM distributes DNS name resolution requests to the first available virtual server in a pool. GTM starts at the top of a manually configured list of virtual servers and sends requests to the first available virtual server in the list. Only when the virtual server becomes unavailable does GTM send requests to the next virtual server in the list. Over time, the first virtual server in the list receives the most requests and the last virtual server in the list receives the least requests.	Use Global Availability when you have specific virtual servers that you want to handle most of the requests.
None	GTM distributes DNS name resolution requests skipping either the next available pool in a multiple pool configuration or the current load balancing method. If all pools are unavailable, GTM returns an aggregate of the IP addresses of all the virtual servers in the pool using BIND.	Use None for the alternate and fallback methods when you want to limit each pool to a single load balancing method. If the preferred load balancing method fails, GTM offers the next pool in a load balancing response.
Ratio	GTM distributes DNS name resolution requests among the virtual servers in a pool or among pools in a multiple pool configuration using weighted round robin, a load balancing pattern in which requests are distributed among several resources based on a priority level or weight assigned to each resource.	Use Ratio when you want to send twice as many connections to a fast server and half as many connections to a slow server.
Return to DNS	GTM immediately distributes DNS name resolution requests to an LDNS for resolution.	Use Return to DNS when you want to temporarily remove a pool from service. You can also use Return to DNS when you want to limit a pool in a single pool configuration to only one or two load balancing attempts.

Round Robin	GTM distributes DNS name resolution requests in a circular and sequential pattern among the virtual servers in a pool. Over time each virtual server receives an equal number of requests.	Use Round Robin when you want to distribute requests equally among all virtual servers in a pool.
Static Persist	GTM distributes DNS name resolution requests to the first available virtual server in a pool using the persist mask with the source IP address of the LDNS and a hash algorithm to determine the order of the virtual servers in the list. This hash algorithm orders the virtual servers in the list differently for each LDNS that is passing traffic to the system taking into account the specified CIDR of the LDNS. Each LDNS (and thus each client) generally resolves to the same virtual server; however, when the selected virtual server becomes unavailable, GTM sends requests to another virtual server until the original virtual server becomes available. Then GTM again resolves requests to that virtual server.	Use Static Persist when you want requests from a specific LDNS to resolve to a specific virtual server.
Topology	GTM distributes DNS name resolution requests using proximity-based load balancing. GTM determines the proximity of the resource by comparing location information derived from the DNS message to the topology records in a topology statement you have configured.	Use Topology when you want to send requests from a client in a particular geographic region to a data center or server located in that region.

Dynamic load balancing methods

This table describes the dynamic load balancing methods available in GTM.

NAME	DESCRIPTION	RECOMMENDED USE
Completion Rate	GTM distributes DNS name resolution requests to the virtual server that currently maintains the least number of dropped or timed-out packets during a transaction between a data center and the client's LDNS.	
CPU	GTM distributes DNS name resolution requests to the virtual server that currently has the most CPU processing time available.	
Hops	GTM distributes DNS name resolution requests to a virtual server in the data center that has the fewest router hops from the client's LDNS. GTM uses the traceroute utility to track the number of router hops between a client's LDNS and each data center.	
Kilobytes per Second	GTM distributes DNS name resolution requests to the virtual server that is currently processing the fewest number of kilobytes per second. Use Kilobytes/Second only with virtual servers for which GTM can collect the kilobytes per second metric.	
Least Connections	GTM distributes DNS name resolution requests to virtual servers on LTM that currently hosts the fewest connections. Use Least Connections only with LTM servers.	
Packet Rate	GTM distributes DNS name resolution requests to the virtual server that is currently processing the fewest number of packets per second.	

Quality of Service	GTM distributes DNS name resolution requests to virtual servers based on a score assigned to each virtual server that is calculated from current performance metrics. Use Quality of Service only when you have configured GTM to calculate an overall score for each virtual server based on performance metrics.	
Round Trip Time	GTM distributes DNS name resolution requests to the virtual server with the fastest measured round trip time between a data center and a client's LDNS.	
Virtual Server Score	GTM distributes DNS name resolution requests to virtual servers on LTM based on a user-defined ranking. Use Virtual Server Score only with LTM systems on which you have assigned scores to each virtual server.	
Virtual Server Capacity	GTM distributes DNS name resolution requests to virtual servers in a list that are weighted by the number of available virtual servers in the pool. The pool with the most available virtual servers is sent more requests; however, over time all the virtual servers in all the pools are sent requests. If more than one virtual server has the same weight, then GTM distributes DNS requests among those virtual servers using the round-robin load balancing method.	

Objective - 1.10 - Identify applicable iRules events including application to Wide-IP versus Listener

1.10 - Identify the purpose and use of Wide-IP

[Link to Online Topic Content](#)

Wide-IP

A Wide-IP is a mapping of a fully qualified domain name (FQDN) to a set of virtual servers that host the domains content, such as a web site, an e-commerce site, or a CDN. Wide-IPs use pools to organize virtual servers, which creates a tiered load balancing effect: the GTM first load balances requests to a Wide-IP to the appropriate pool, then load balances within the pool to the appropriate virtual server.

Wide-IPs also support iRules for further managing and directing network traffic. An iRule is a set of one or more TCL-based expressions that direct network traffic beyond load balancing operations.

A Wide-IP does not require iRules to operate effectively. However, iRules are a powerful mechanism for customizing how the GTM handles network connection requests.

1.10 - Identify LTM iRule events versus GTM iRule events (Apply to Wide-IP vs Listener)

[Link to Online Topic Content](#)

Introduction

As you work with GTM, you might find that you want to incorporate additional customizations beyond the available features associated with load balancing, monitors, or other aspects of your traffic management. For example, you might want to have the system respond to a name resolution request with a specific CNAME record, but only when the request is for a particular Wide-IP and originates from Europe. In GTM, these customizations are defined through iRules®. iRules are code snippets that are based on TCL 8.4. These snippets allow you a great deal of flexibility in managing your global network traffic.

If you are familiar with Local Traffic Manager, you might already be aware of and use iRules to manage your network traffic on a local level. The iRules in GTM share a similar syntax with their Local Traffic Manager counterparts, but support a different set of events and objects.

Due to the dynamic nature of iRules development, the following sections focus on providing an overview of iRule operations and describe the events and command specific to GTM. For additional information about how to write iRules, visit the F5 DevCentral web site: <http://devcentral.f5.com>. At this site, you can learn more about iRules development, as well as discuss iRules functionality with others.

What is an iRule?

An iRule is a script that you write if you want individual connections to target a pool other than the default pool defined for a virtual server. iRules allow you to more directly specify the pools to which you want traffic to be directed. Using iRules, you can send traffic not only to pools, but also to individual pool members or hosts.

The iRules you create can be simple or sophisticated, depending on your content-switching needs.

Example of an iRule

```
when DNS_REQUEST {  
    if { [IP::addr [IP::client_addr] equals 10.10.10.10] } {  
        pool my_pool  
    }  
}
```

This iRule is triggered when a DNS request has been detected, causing GTM to send the packet to the pool `my_pool`, if the IP address of the local DNS making the request matches `10.10.10.10`.

iRules can direct traffic not only to specific pools, but also to individual pool members, including port numbers and URI paths, either to implement persistence or to meet specific load balancing requirements.

The syntax that you use to write iRules is based on the Tool Command Language (TCL) programming standard. Thus, you can use many of the standard TCL commands, plus a set of extensions that GTM provides to help you further increase load balancing efficiency.

Within GTM, you assign iRules to the Wide-IPs in your network configuration.

Event-based traffic management

In a basic system configuration where no iRule exists, GTM directs incoming traffic to the default pool assigned to the Wide-IP that receives that traffic based on the assigned load balancing modes. However, you might want GTM to direct certain kinds of connections to other destinations. The way to do this is to write an iRule that directs traffic to that other destinations contingent on a certain type of event occurring. Otherwise, traffic continues to go to the default pool assigned to the Wide-IP.

iRules are evaluated whenever an event occurs that you have specified in the iRule. For example, if an iRule includes the event declaration `DNS_REQUEST`, then the iRule is triggered whenever GTM receives a name resolution request. GTM then follows the directions in the remainder of the iRule to determine the destination of the packet.

When you assign multiple iRules as resources for a Wide-IP, it is important to consider the order in which you list them on the Wide-IP. This is because GTM processes duplicate iRule events in the order that the applicable iRules are listed. An iRule event can therefore terminate the triggering of events, thus preventing GTM from triggering subsequent events.

Event declarations

The iRules feature includes several types of event declarations that you can make in an iRule. Specifying an event declaration determines when GTM evaluates the iRule. The following sections list and describe these event types.

iRule Event	Description
DNS_REQUEST	Triggered when a DNS request is received from a client.
LB_SELECTED	Triggered when the Global Traffic Manager has selected a target node.
LB_FAILED	Triggered when a connection to the server was unable to complete. This might occur if the pool has no available members or a selected pool member is otherwise not available.
RULE_INIT	Triggered when an iRule that contains the RULE_INIT event is changed, or when the gtm utility restarts. Note that only the following commands are valid with this event: whoami, whereami, crc32, findstr, log, substr, and whereis.

Objective - 1.11 - Identify the purpose of GTM tools and when to use them [checkcert, iqdump, etc.]

1.11 - Identify the purpose of GTM tools and when to use them [checkcert, iqdump, etc.]

[Link to Online Topic Content](#)

GTM Tools

The troubleshooting tools checkcert and iqdump are available from the command line interface on the GTM platform.

checkcert

The checkcert utility examines all the certificates in the /config/ssl/ssl.crt directory, including bundled certificates files. The checkcert utility is called from /etc/cron.weekly/5checkcert, and it examines the expiration date of each certificate on the system.

Note: By default, the checkcert utility does not check the ca-bundle.crt certificate file.

If the checkcert utility finds a certificate that has expired or will expire within 30 days, it logs an error message to the /var/log/ltn file.

[Link to Online Topic Content](#)

iqdump

The command iqdump can be used to view the data transmitted between systems using iQuery. If the remote LTM devices are marked green by the GTM systems, but the GTM systems fail to automatically discover the virtual server and link objects, you can use the iqdump utility to verify that the LTM is properly broadcasting the virtual server and link objects.

GTM Configuration Scripts

The GTM configuration scripts allow you to establish communications between the GTM systems and other external BIG-IP systems. Before you run any of the GTM configuration scripts, you should use the GTM Configuration utility to define any remote

BIG-IP systems (including other GTM systems), with which GTM will communicate. Once you have defined the other systems with which the GTM system will communicate, you may need to run one or more of the GTM configuration scripts. The GTM configuration scripts are defined as follows:

[Link to Online Topic Content](#)

big3d_install

The big3d process runs on all BIG-IP systems, and provides metrics collection data for BIG-IP systems. The big3d_install script is an interactive script that allows you to install the current version of the big3d process on remote F5 systems. If the current or newer version of the big3d process is found to be running on the remote BIG-IP system, installation is skipped for that BIG-IP system. The big3d_install script also copies the trusted device certificate from the local GTM system to the /config/big3d/client.crt file on the remote BIG-IP system, and the trusted server certificate from the remote BIG-IP system to the /config/gtm/server.crt file on the local GTM system.

[Link to Online Topic Content](#)

bigip_add

The bigip_add script is an interactive script that exchanges iQuery SSL certificates with a remote BIG-IP system. The bigip_add script appends the local GTM system's SSL certificate to the remote BIG-IP system's list of authorized certificates (contained in the /config/big3d/client.crt file). The script then appends the remote BIG-IP system's iQuery SSL certificate to the GTM system's local list of authenticated iQuery SSL certificates (/config/gtm/server.crt).

[Link to Online Topic Content](#)

gtm_add

The gtm_add script can be used to integrate a new GTM system into an existing sync group that has one or more remote GTM Controllers. The script will replace the current configuration files (bigip_gtm.conf, named.conf and named zone files) on GTM on which it is run with the configuration of the remote GTM system in the specified sync group.

Objective - 1.12 - Explain how zone transfers work [multi master, master/slave, DNSExpress, incremental/full, updates (notify/expire)]

1.12 - Explain how zone transfers work [multi master, master/slave, DNSExpress, incremental/full, updates (notify/expire)]

[Link to Online Topic Content](#)

Ref: 1, pp. 1024.

Zone Transfer

A zone transfer is a TCP-based client–server request. The client requesting a zone transfer may be a primary server (master) or a secondary server (slave), requesting data from a primary server of the zone. The portion of the database that is replicated is a zone.

By default, GTM is configured to secure BIND to not allow zone transfers except from the localhost. However, you can configure GTM to allow zone file transfers to other DNS servers.

Zone transfer comprises a preamble followed by the actual data transfer. The preamble comprises a lookup of the SOA (Start of Authority) resource record for the "zone apex", the node of the DNS namespace that is at the top of the "zone". The fields of this SOA resource record, in particular the "serial number", determine whether the actual data transfer need occur at all. The client compares the serial number of the SOA resource record with the serial number in the last copy of that resource record that it has. If the serial number of the record being transferred is greater, the data in the zone are deemed to have "changed" (in some fashion) and the slave proceeds to request the actual zone data transfer. If the serial numbers are identical, the data in the zone are deemed not to have "changed", and the client may continue to use the copy of the database that it already has, if it has one.

The actual data transfer process begins by the client sending a query (opcode 0) with the special QTYPE (query type) AXFR (value 252) over the TCP connection to the server. The server responds with a series of response messages, comprising all of the resource records for every domain name in the "zone". The first response comprises the SOA resource record for the zone apex. The other data follows in no specified order. The

end of the data is signaled by the server repeating the response containing the SOA resource record for the zone apex.

Some zone transfer clients perform the SOA lookup of the preamble using their system's normal DNS query resolution mechanism. These clients do not open a TCP connection to the server until they have determined that they need to perform the actual data transfer. However, since TCP can be used for normal DNS transactions, as well as for zone transfer, other zone transfer clients perform the SOA lookup preamble over the same TCP connection as they then (may) perform the actual data transfer. These clients open the TCP connection to the server before they even perform the preamble.

The preceding describes full zone transfer. Incremental zone transfer differs from full zone transfer in the following respects:

- The client uses the special QTYPE IXFR (value 251) instead of the AXFR QTYPE.
- The client sends the SOA resource record for the zone apex that it currently has, if any, in the IXFR message, letting the server know which version of the "zone" it believes to be current.
- Though the server may respond in the normal AXFR manner with the full data for the zone, it may also instead respond with an "incremental" data transfer. This latter comprises the list of changes to the zone data, in zone serial number order, between the version of the zone that the client reported to the server as having and the version of the zone that is current at the server. The changes comprise two lists, one of resource records that are deleted and one of resource records that are inserted. (A modification to a resource record is represented as a deletion followed by an insertion.)

Zone transfer is entirely client-initiated. Though servers can send a NOTIFY message to clients (that they have been informed about) whenever a change to the zone data has been made, the scheduling of zone transfers is entirely under the control of the clients. Clients schedule zone transfers initially, when their databases are empty, and thereafter at regular intervals, in a pattern controlled by the values in the "refresh", "retry", and "expire" fields in the SOA resource record of the zone apex.

[Link to Online Topic Content](#)

DNSExpress

DNSExpress allows you to transfer DNS zones from your current infrastructure to the BIG-IP. The BIG-IP can then answer requests for those zones. DNSExpress doesn't run full BIND, so it's not as vulnerable as a typical BIND infrastructure. DNS Express provides the ability for a BIG-IP to act as a high speed, authoritative secondary DNS server. This allows the BIG-IP to perform zone transfers from multiple primary DNS servers that are responsible for different zones, perform a zone transfer from the local BIND server on the BIG-IP, and serve DNS records faster than the primary DNS servers and the local BIND server.

To use DNS Express, you need to create a DNS Express zone. Then, you can transfer zone records from the local BIND server or back-end DNS servers to DNS Express. Note that DNS Express is configured under "Local Traffic" as part of the Local Traffic Manager (LTM).

Objective - 1.13 - Given a scenario determine the impact of a custom DNS profile for various types of queries, determine what response will be given and where it will come from

1.13 - Given a scenario determine the impact of a custom DNS profile for various types of queries; determine what response will be given and where it will come from

Scenario Based Questions

This topic is focused on the different settings within the DNS profile. Knowing what the different settings do and how they affect name resolutions will give the candidate the ability to answer the questions on this topic.

To prepare for scenario based questions the candidate will need to complete hands-on configuration and testing of the configuration on the GTM. This will allow the candidate to better understand how different configurations can produce different results. All F5 exams use scenario-based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

1.13 - Explain all of the features that can be enabled in a DNS profile (DNS cache, unhandled query, DNS Express, enable GTM, enable bind)

[Link to Online Topic Content](#)

DNS Profiles

You can create a custom DNS profile to enable various features such as converting IPv6-formatted addresses to IPv4 format, enabling DNS Express, and enabling DNSSEC.

General property	Description	Default Value
Name	Specifies the user-supplied name of the profile. Specifying a name for your profile is required.	No default value
Parent Profile	Specifies the profile from which your custom profile is derived.	dns
Global Traffic Management	Specifies whether the system uses Global Traffic Manager to manage the response.	Enabled
DNS IPv6 to IPv4	Specifies whether you want the BIG-IP system to convert IPv6-formatted IP addresses to IPv4-formatted IP addresses. The possible values are: Disabled: Indicates that the BIG-IP system does not map IPv4 addresses to IPv6 addresses. Secondary: Indicates that the BIG-IP system receives an AAAA query and forwards the query to a DNS server. Immediate: Indicates that the BIG-IP system receives an AAAA query and forwards the query to a DNS server. v4 Only: Indicates that the BIG-IP system receives an AAAA query, but forwards an A query to a DNS server. When you select Secondary, Immediate, or v4 Only, you must also provide a prefix in the IPv6 to IPv4 Prefix field and make a selection from the IPv6 to IPv4 Additional Section Rewrite list.	Disabled
DNS Express	Indicates whether the dns-express service is enabled. The service handles zone transfers from the primary DNS server.	Enabled
DNSSEC	Specifies whether the system signs responses and replies to DNSSEC-specific queries (for example, DNSKEY query type).	Enabled

Unhandled Query Actions	<p>Specifies whether the system uses the local BIND server on the BIG-IP system. The possible values are:</p> <p>Allow: Indicates that the BIG-IP system forwards the connection request to another DNS server or DNS server pool. If a DNS server pool is not associated with a listener and the Use BIND Server on BIG-IP setting is set to Enabled, connection requests are forwarded to the local BIND server.</p> <p>Drop: Indicates that the BIG-IP system does not respond to the query.</p> <p>Reject: Indicates that the BIG-IP system returns the query with the REFUSED return code.</p> <p>Hint: Indicates that the BIG-IP system returns the query with a list of root name servers.</p> <p>No Error: Indicates that the BIG-IP system returns the query with the NOERROR return code.</p>	Allow
Use BIND Server on BIG-IP	<p>Specifies whether the system forwards non-Wide-IP queries to the local BIND server on the BIG-IP system. For best performance, disable this setting when using a DNS cache.</p>	Enabled
Process Recursion Desired	<p>Indicates whether to process client-side DNS packets with Recursion Desired set in the header. If set to Disabled, processing of the packet is subject to the unhandled-query-action option.</p>	Enabled
DNS Cache	<p>Specifies whether the system caches DNS responses.</p> <p>Enabled: Indicates the BIG-IP system caches DNS responses handled by the virtual servers associated with this profile. When you enable this setting, you must also select a cache from the DNS Cache Name list.</p> <p>Disabled: Indicates the BIG-IP system does not cache DNS responses handled by the virtual servers associated with this profile. However, the profile retains the association with the DNS cache in the DNS Cache Name field. Disable this setting when you want to debug the system.</p>	Disabled
DNS Cache Name	<p>Specifies the user-created cache that the system uses to cache DNS responses. When you select a cache for the system to use, you must also enable the DNS Cache setting.</p>	No default value
DNS Security	<p>Indicates whether DNS firewall capability is enabled.</p>	Disabled
DNS Security Profile Name	<p>Specifies the DNS security profile to use.</p>	No default value

Logging	Indicates whether to enable high speed logging for DNS queries and responses. When it is set to Enabled, you must also specify a Logging Profile.	Disabled
Logging Profile	Specifies the DNS logging profile used to configure what events get logged and their message format. These are the DNS Logging profiles you create using the Other option under Profiles in the BIG-IP Configuration utility.	No default value

Objective - 1.14 - Given a scenario with a specific query source IP address and various pool and Wide-IP loading balancing methods and topology rules/regions determine the response that will be given

1.14 - Given a scenario with a specific query source IP address and various pool and Wide-IP loading balancing methods and topology rules/regions determine the response that will be given

[Link to Online Topic Content](#)

Scenario Based Questions

This topic is focused on the results of a Wide-IP name query and the different ways the Wide-IP will resolve for the name, with configurations using different load balancing methods. Experience with configuring a Wide-IP and testing resolutions against it will give the candidate the ability to answer the questions on this topic.

To prepare for scenario based questions the candidate will need to complete hands-on configuration and testing of the configuration on the GTM. This will allow the candidate to better understand how different configurations can produce different results. All F5 exams use scenario-based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

Objective - 1.15 - Explain sync group/iQuery purpose, configuration and basic requirements

1.15 - Explain how iQuery is used in sync groups and LTM monitoring

[Link to Online Topic Content](#)

Sync Groups/iQuery

A GTM synchronization group is a collection of multiple GTM systems that share and synchronize configuration settings. You must meet several minimum requirements for GTM synchronization group members to communicate and synchronize properly.

[Link to Online Topic Content](#)

For the GTM synchronization group members to properly synchronize their configuration settings, verify that the following requirements are in place:

- **GTM synchronization group members must be running the same software version**

A GTM device should be running the same software version as other members in the synchronization group. GTM devices that are running different software versions will not be able to communicate and properly synchronize GTM configuration and zone files. For information about displaying the software version, refer to SOL8759: Displaying the BIG-IP software version.

- **Synchronization parameters must be properly defined for all members**

Synchronization must be enabled and each device must have the same synchronization group name. You can define the synchronization parameters by navigating to:

GTM 10.0.0 - 11.4.1:

System > Configuration > Device > GTM > General

- **NTP must be configured on each device**

Before you can synchronize GTM systems, you must define the Network Time Protocol (NTP) servers for all synchronization group members. Configuring NTP servers ensures that each GTM synchronization group member is referencing the same time when verifying the configuration data that needs to be synchronized. You can configure NTP by navigating to System > Configuration > Device > NTP.

- **Port Lockdown must be set properly for the relevant self-IP addresses**

Port lockdown is a security feature that specifies the protocols and services from which a self-IP address can accept traffic. F5 recommends using the Allow Default option for self-IP addresses that are used for synchronization and other critical redundant pair intercommunications. You can configure port lockdown by navigating to Network > Self IPs.

- **TCP port 4353 must be allowed between GTM systems**

GTM synchronization group members use TCP port 4353 to communicate. You must verify that port 4353 is allowed between GTM systems.

- **Compatible big3d versions must be installed on synchronization group members**

The big3d process runs on BIG-IP systems and collects performance information on behalf of the GTM system. For metrics collection to work properly, synchronization group members must run the same version of the big3d process. For more information about verifying big3d version information, refer to SOL13703: Overview of big3d version management.

- **A valid device certificate must be installed on all members**

The device certificate is used by the F5 system to identify itself to a requesting F5 client system. The default device certificate, `/config/httpd/conf/ssl.crt/server.crt`, must be installed on each sync group member. You can verify the certificate validity by navigating to System > Device Certificates.

Objective - 1.16 - Explain the networking requirements of placing devices within a GTM data center object

1.16 - Explain and identify GTM objects (Data center, link, server, virtual server, prober pool, pool, Wide-IP)

[Link to Online Topic Content](#)

GTM Configuration Objects

The following object types are found within a GTM configuration. GTM consolidates the paths and metrics data collected from the servers, virtual servers, and links in the data center, and uses that data to conduct load balancing and route DNS name resolution requests to the best-performing site based on different factors.

Data Centers

The GTM considers that all of the resources in your networks are located in a data center. The Data Center objects in the GTM configuration represent the physical locations for all of your network's equipment.

Links

A link is a logical representation of a physical device (router) that connects your network to the Internet. GTM tracks the performance of links, which influence the availability of pools, data centers, Wide-IPs, and distributed applications.

Prober Pool

A Prober pool is an ordered collection of one or more BIG-IP systems. GTM can be a member of more than one Prober pool, and a Prober pool can be assigned to an individual server or a data center. When you assign a Prober pool to a data center, by default, the servers in that data center inherit that Prober pool.

The members of a Prober pool perform monitor probes of servers to gather data about the health and performance of the resources on the servers. GTM makes load balancing decisions based on the gathered data. If all of the members of a Prober pool are marked down, or if a server has no Prober pool assigned, GTM reverts to a default intelligent probing algorithm to gather data about the resources on the server.

Servers

A server defines a physical system on the network. Servers contain the virtual servers that are the ultimate destinations of DNS name resolution requests. GTM supports three types of servers:

- BIG-IP systems
 - Any member of the BIG-IP system product line.
- Third-party load balancing systems
 - A third-party load balancing system is any system, other than a BIG-IP system, that supports and manages virtual servers on the network.
- Third-party host servers
 - A third-party host server is any server on the network that does not support virtual servers.

Important: At a minimum, you must define two servers, one that represents GTM and one that represents another managed server (either a load balancing or host server).

Virtual Servers

A virtual server is a specific IP address and port number that points to a resource on the network. In the case of host servers, this IP address and port number likely point to the resource itself. With load balancing systems, such as the Local Traffic Manager™, virtual servers are often proxies that allow the load balancing server to manage a resource request across a multitude of resources.

Pool

A pool is a collection of virtual servers that can reside on multiple servers. A virtual server is a combination of IP address and port number that points to a specific resource on the network. When you add a virtual server to a pool, it becomes a pool member. A pool member is a virtual server that has attributes that pertain to the virtual server only in the context of the pool. A virtual server can be a member of multiple pools and have different attributes in each pool. GTM directs traffic to a pool member, based on the attributes of the pool member.

Wide-IP

A Wide-IP maps a fully qualified domain name (FQDN) to one or more pools of virtual servers that host the content of a domain. When an LDNS issues a DNS name resolution for a Wide-IP, the configuration of the Wide-IP indicates which pools of virtual servers are eligible to respond to the request, and which load balancing methods GTM uses to select the pool.

Section 2 - Deployment

Objective - 2.01 Explain when to configure translation addresses for local data center connectivity

2.01 - Explain when to configure translation addresses for local data center connectivity

[Link to Online Topic Content](#)

Translation Addresses

A translation address is used on a configuration object's IP address in a GTM configuration when there is a NAT performed on that IP (by Firewall or other device).

When a GTM system sends iQuery probes to a BIG-IP device that resides in an infrastructure that uses firewall NAT rules, you must configure the GTM to use the firewall NAT addresses (translation address) of the BIG-IP device as the destination addresses for the probes. Additionally, you must configure the GTM system to recognize what the configured private network self-IP and Virtual addresses are when the system analyzes the iQuery metrics that the iQuery probes return. When you add a BIG-IP device, as a server, to the GTM system, the address of the device should be the firewall NAT address, and the private network address that is configured on the BIG-IP device is the stated translation address for that server.

Objective - 2.02 Explain how to configure GTM sync groups and iQuery

2.02 - Explain how to configure GTM sync groups and iQuery

[Link to Online Topic Content](#)

Ref: 3, pp. 9-14.

Configuring Sync Groups

Each Global Traffic Manager that you synchronize must belong to a specific group of systems, called a synchronization group. A synchronization group is a collection of multiple Global Traffic Manager systems that share and synchronize configuration settings. Initially, when you enable synchronization for a Global Traffic Manager, the system belongs to a synchronization group called default. However, you can create new groups at any time to customize the synchronization process, ensuring that only certain sets of Global Traffic Manager systems share configuration values.

To create a synchronization group

1. On the Main tab of the navigation pane, expand System and then click Configuration.

The general properties screen opens.

2. From the Global Traffic menu, choose General.

The general global properties screen opens.

3. In the Synchronization Group Name box, type a name of either an existing synchronization group, or a new group.

Note: When you change the name of a synchronization group, the new name is synchronized to all systems that belong to that synchronization group.

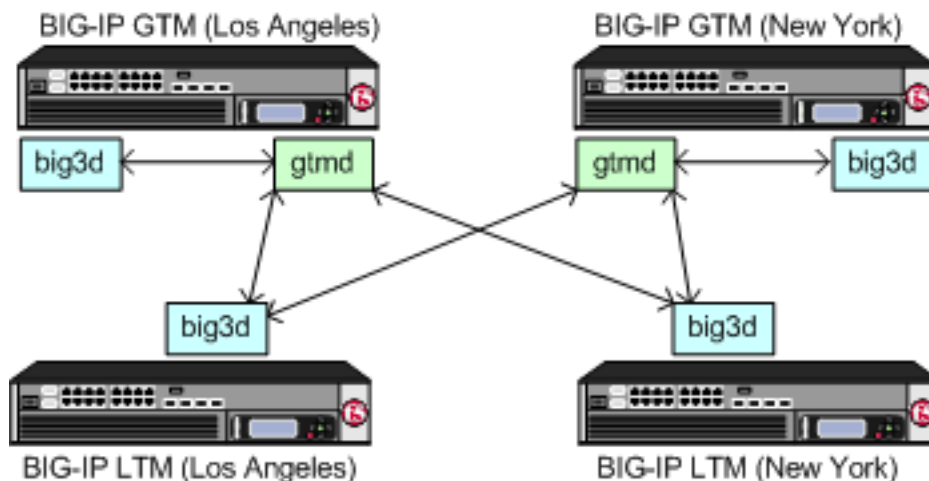
4. Click the Update button to save your changes.

About iQuery and communications between BIG-IP systems

When the GTM communicates with other BIG-IP systems, such as LTM systems, it uses a proprietary protocol called iQuery to send and receive information. If the GTM is communicating with another BIG-IP system, it uses the big3d utility to handle the communication traffic. If the GTM is instead communicating with another GTM, it uses a different utility, called gtmd, which is designed for that purpose.

Part of the process when establishing communications between the GTM and other BIG-IP systems is to open port 22 and port 4353 between the two systems. Port 22 allows the GTM to copy the newest version of the big3d utility to existing systems, while iQuery requires the port 4353 for its normal communications.

In order for other BIG-IP systems to communicate with the GTM, F5 Networks recommends that you update the big3d utility on older BIG-IP systems by running the big3d_install script from the GTM.



Objective - 2.03 Given a set of requirements select the appropriate load balancing methods [ex. Wide-IP level, pool level, different types and combinations]

2.03 - Given a scenario determine the load balancing decision based on virtual server status and configure load balancing (single pool versus multiple pools, effect of secondary and fallback mechanisms in the first pool, effect of topology and topology records at the Wide-IP level versus pool level, iRule effects)

[Link to Online Topic Content](#)

Scenario Based Questions

This topic is focused on the results of a Wide-IP name query and the different levels of load balancing that can occur within the different settings inside the configuration of the Wide-IP. Understanding how a Wide-IP can load balance across multiple pools with multiple methods of load balancing, and once a pool has been chosen load balancing across the members of the pool will occur to determine the IP to resolve based on multiple methods of load balancing. Experience with configuring a Wide-IP and testing resolutions against it will give the candidate the ability to answer the questions on this topic.

To prepare for scenario based questions the candidate will need to complete hands-on configuration and testing of the configuration on the GTM. This will allow the candidate to better understand how different configurations can produce different results. All F5 exams use scenario-based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

Understanding Load Balancing on the GTM

When the GTM receives a name resolution request, the system employs a load balancing mode to determine the best available virtual server to which to send the request. Once the GTM identifies the virtual server, it constructs a DNS answer and sends that answer back to the requesting clients local DNS server. The DNS answer, or resource record, can be either an A record that contains the IP address of the virtual server, or a CNAME record that contains the canonical name for a DNS zone.

Within the GTM, there are two categories of load balancing modes from which to select: static and dynamic. A static load balancing mode selects a virtual server based on a pre-

defined pattern. A dynamic load balancing mode selects a virtual server based on current performance metrics.

The GTM provides a tiered load balancing system. A tiered load balancing system is a load balancing system that occurs at more than one point during the resolution process. The tiers within the GTM are as follows:

Wide IP-level load balancing

A wide IP contains two or more pools. The GTM load balances requests, first to a specific pool, and then to a specific virtual server in the selected pool. If the preferred, alternate, and fallback load balancing methods that are configured for the pool or virtual server fail, then the requests fail, or the system falls back to DNS.

Pool-level load balancing

A pool contains one or more virtual servers. After the GTM uses wide IP-level load balancing to select the best available pool, it uses a pool-level load balancing to select a virtual server within that pool. If the first virtual server within the pool is unavailable, the GTM selects the next best virtual server based on the load balancing mode assigned to that pool.

For each pool that you manage, the GTM supports three types of load balancing methods: preferred, alternate, and fallback. The preferred load balancing method is the load balancing mode that the system attempts to use first. If the preferred method fails to provide a valid resource, the system uses the alternate load balancing method. Should the alternate load balancing method also fail to provide a valid resource, the system uses the fallback method.

One of the key differences between the alternate methods and the other two load balancing methods is that only static load balancing modes are available from the alternate load balancing list. This limitation exists because dynamic load balancing modes, by definition, rely on metrics collected from different resources. If the preferred load balancing mode does not return a valid resource, it is likely that the GTM was unable to acquire the proper metrics to perform the load balancing operation. By limiting the alternate load balancing options to static methods only, the GTM can better ensure that, should the preferred method prove unsuccessful, the alternate method returns a valid result.

Objective - 2.04 Given a scenario select the appropriate deployment type: screening mode, DNS delegation, caching resolver, and DNS 6 to 4

[Link to Online Topic Content](#)

Scenario Based Questions

This topic is focused on the different types of GTM deployment options. Experience with configuring each of these deployment types will give the candidate the ability to answer the questions on this topic.

To prepare for scenario based questions the candidate will need to complete hands-on configuration and testing of the configuration on the GTM. This will allow the candidate to better understand how different configurations can produce different results. All F5 exams use scenario based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

GTM Deployment types

Screening Mode

The Authoritative Screening architecture enables the GTM to receive all DNS queries, managing very high-volume DNS by load balancing requests to a pool of DNS servers. Additionally, the Authoritative Screening architecture seamlessly provides all of the benefits of intelligent GSLB services.

When a DNS query is received, the BIG-IP checks the record type. If the type is an A, AAAA, A6, or CNAME request, it is sent to GTM module. The GTM checks each request and response, looking for a match against the Wide-IP list of FQDN names. If there is a match, the GTM performs the appropriate GSLB functions and return the best IP address appropriate for the requesting client.

If the DNS request does not match the Wide-IP list, GTM passes the request to a pool of DNS servers, which provides an additional layer of scalability and availability, increasing the query performance and ensuring optimal uptime of DNS services. Screening mode simplifies management when used with other DNS servers.

GTM inspects all DNS responses from the DNS servers. If the response contains a DNS name that matches a Wide-IP, GTM intercepts the response, applies the GTM operations for that item, and re-writes the response before sending it on to the client.

DNS Delegation

When a client requests DNS resolution for the host name `www.domain.com`, the DNS server that is authoritative for the `domain.com` domain responds with the CNAME (alias) record `www.wip.domain.com`. The client then requests resolution for the host name `www.wip.domain.com`. Since the `domain.com` zone points to `gtm1.domain.com` and `gtm2.domain.com` as the authoritative name servers for the subdomain `wip.domain.com`, the resolution request for `www.wip.domain.com` is then sent to one of the GTM systems. The GTM system responds with the most appropriate A (address) record based on the Wide-IP configuration.

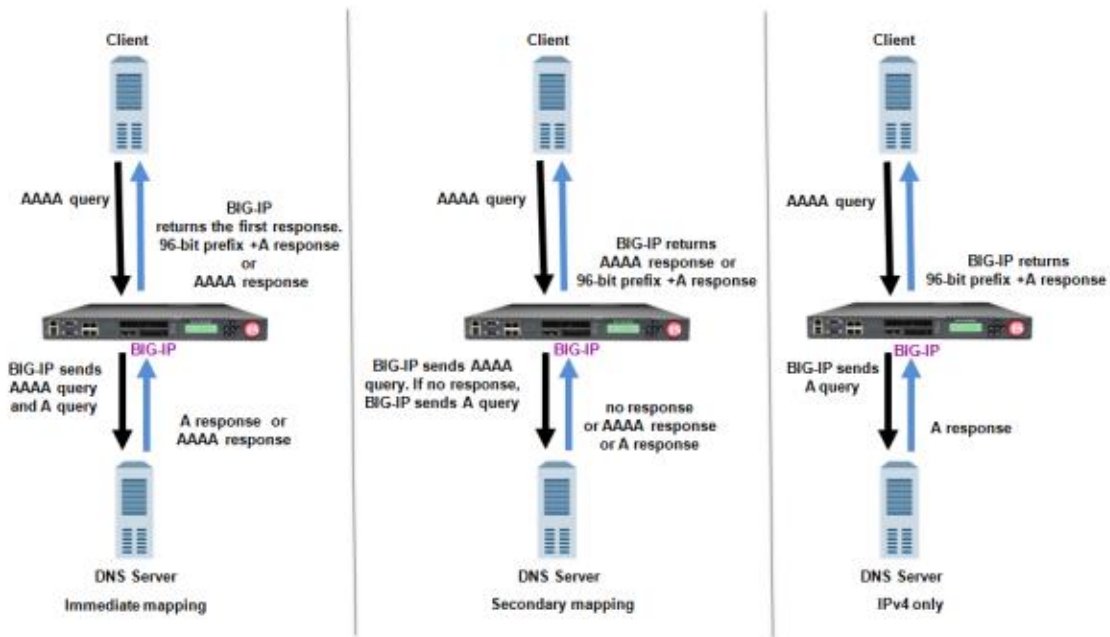
Caching Resolver

You can configure a resolver cache on the BIG-IP system to resolve DNS queries and cache the responses. The next time the system receives a query for a response that exists in the cache, the system returns the response from the cache. The resolver cache contains messages, resource records, and the name servers that the system queries to resolve DNS queries.

It is important for network architects to note that it is possible to configure the local BIND instance on the BIG-IP system to act as an external DNS resolver. However, F5 Networks does not recommend this approach, because the performance of BIND is slower than using a resolver cache.

DNS64

You can configure LTM and GTM systems to handle IPv6-only client connection requests to IPv4-only servers on your network by returning an AAAA record response to the client.



Mapping IPv6 addresses to IPv4 addresses

2.04 - Determine when to use ZoneRunner to manage DNS records on GTM

[Link to Online Topic Content](#)

ZoneRunner

ZoneRunner is a graphical GUI for editing BIND configuration files including named.conf and creating database files for individual zones. If you want to manage your zone files using the GTM, you can use the ZoneRunner utility to create and manage DNS zone files and configure the BIND instance on GTM.

With the ZoneRunner utility, you can:

- Import and transfer DNS zone files
- Manage zone resource records
- Manage views
- Manage a local name server and the associated configuration file, named.conf
- Transfer zone files to a name server
- Import only primary zone files from a name server

Objective - 2.05 Explain how to configure GTM to return non-Wide-IP supported records [ex. MX, SRV, TXT records, etc.]

2.05 - Explain how to configure GTM to return non-Wide-IP supported records [ex. MX, SRV, TXT records, etc.]

[Link to Online Topic Content](#)

Returning Non-Wide-IP Supported Records

For GTM to return non-Wide-IP supported records you must have ZoneRunner configured. ZoneRunner is the GUI for Bind on the system, which supports the following DNS file types:

DNS FILE TYPE	DESCRIPTION
SOA (Start of authority)	The start of authority resource record, SOA, starts every zone file and indicates that a name server is the best source of information for a particular zone. The SOA record indicates that a name server is authoritative for a zone. There must be exactly one SOA record per zone. Unlike other resource records, you create a SOA record only when you create a new master zone file.
A (Address)	The Address record, or A record, lists the IP address for a given host name. The name field is the host's name, and the address is the network interface address. There should be one A record for each IP address of the machine.
AAAA (IPv6 Address)	The IPv6 Address record, or AAAA record, lists the 128-bit IPv6 address for a given host name.
CNAME (Canonical Name)	The Canonical Name resource record, CNAME, specifies an alias or nickname for the official, or canonical, host name. This record must be the only one associated with the alias name. It is usually easier to supply one A record for a given address and use CNAME records to define alias host names for that address.

DNAME (Delegation of Reverse Name)	The Delegation of Reverse Name resource record, DNAME, specifies the reverse lookup of an IPv6 address. These records substitute the suffix of one domain name with another. The DNAME record instructs Global Traffic Manager (or any DNS server) to build an alias that substitutes a portion of the requested IP address with the data stored in the DNAME record.
HINFO (Host Information)	The Host Information resource record, HINFO, contains information on the hardware and operating system relevant to Global Traffic Manager (or other DNS).
MX (Mail Exchanger)	The Mail Exchange resource record, MX, defines the mail system(s) for a given domain.
NS (name server)	The name server resource record, NS, defines the name servers for a given domain, creating a delegation point and a subzone. The first name field specifies the zone that is served by the name server that is specified in the name servers name field. Every zone needs at least one name server.
PTR (Pointer)	A name pointer resource record, PTR, associates a host name with a given IP address. These records are used for reverse name lookups.
SRV (Service)	The Service resource record, SRV, is a pointer that provides for an alias for a given service to be redirected to another domain. For example, if the fictional company SiteRequest had an FTP archive hosted on archive.siterequest.com, the IT department can create an SRV record that allows an alias, ftp.siterequest.com to be redirected to archive.siterequest.com.
TXT (Text)	The Text resource record, TXT, allows you to supply any string of information, such as the location of a server or any other relevant information that you want available.

Objective - 2.06 Given a scenario of specific virtual server status, pool and Wide-IP load balancing settings determine the answer returned [Single pool versus multiple pools, effect of secondary and fall-back mechanisms in the first pool, effect of topology and topology records at the Wide-IP level versus pool level, and iRule effects]

2.06 - Given a scenario of specific virtual server status, pool and Wide-IP load balancing settings determine the answer returned [Single pool versus multiple pools, effect of secondary and fall-back mechanisms in the first pool, effect of topology and topology records at the Wide-IP level versus pool level, and iRule effects]

[Link to Online Topic Content](#)

Scenario Based Questions

This topic is very similar to topic 2.03.

This topic is focused on the results of a Wide-IP name query and the different levels of load balancing that can occur within the different settings inside the configuration of the Wide-IP. Understanding how a Wide-IP can load balance across multiple pools with multiple methods of load balancing, and once a pool has been chosen load balancing across the members of the pool will occur to determine the IP to resolve based on multiple methods of load balancing including fall-back mechanisms. Experience with configuring a Wide-IP and testing resolutions against it will give the candidate the ability to answer the questions on this topic.

To prepare for scenario based questions the candidate will need to complete hands-on configuration and testing of the configuration on the GTM. This will allow the candidate to better understand how different configurations can produce different results. All F5 exams use scenario-based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

Objective - 2.07 Given a set of topology requirements configure a deployment using user defined topology prefixes

2.07 - Given these topology regions and these rules with load balancing configured as such, what would be the response provided

[Link to Online Topic Content](#)

Scenario Based Questions

This topic is focused on the results of using topology as the load balancing method for a Wide-IP. Experience with configuring a topology record for a Wide-IP and testing resolutions against it will give the candidate the ability to answer the questions on this topic.

To prepare for scenario based questions the candidate will need to complete hands-on configuration and testing of the configuration on the GTM. This will allow the candidate to better understand how different configurations can produce different results. All F5 exams use scenario-based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

Topology Records

Topology is a proximity-based load balancing mode that allows you to direct traffic by defining topology records, and selecting the Topology load-balancing mode for the Wide-IP or pool. The Topology mode bases the distribution of requests on the topology records and the weighted scores configured for each record. The topology records direct DNS queries to the closest virtual server, based on geographical information.

There is only one topology record list for the GTM system, so all Wide-IPs using the topology method on the GTM share the same one. The GTM system looks up topology records in the order they appear in the Configuration utility and the configuration file. As a result, you should place more specific topology records toward the top of the topology statement, and less specific records toward the end of the topology statement. You can change the order of existing topology records by using the Change Order button on the Topology Records page.

2.07 - What is the effect of weighting on topology records?

[Link to Online Topic Content](#)

Weight (score)

The weight specifies the score that will be given to a destination object, which matches the topology record. In the event that a name resolution request matches more than one topology record, the GTM system uses the destination object with the highest weight to determine which statement it uses to load balance the request.



Objective - 2.08 Given a scenario configure a deployment using auto-discovery [behavior of delete versus no-delete with auto-discovery, compatibility with translation, and route domains]

2.08 - Given a scenario configure a deployment using auto-discovery [behavior of delete versus no-delete with auto-discovery, compatibility with translation, and route domains]

[Link to Online Topic Content](#)

Scenario Based Questions

This topic is focused on the using auto-discovery function to automatically build out the virtual servers that are tied to a Server object within a Data Center. Experience with using auto-discovery will give the candidate the ability to answer the questions on this topic.

To prepare for scenario based questions the candidate will need to complete hands-on configuration and testing of the configuration on the GTM. This will allow the candidate to better understand how different configurations can produce different results. All F5 exams use scenario-based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

Auto-discovery

Auto-Discovery is a process by which the GTM system automatically discovers virtual servers and links that are associated with defined BIG-IP systems. When you first enable virtual server or link discovery for a BIG-IP system, the GTM system searches for the resources on the target BIG-IP system, and adds them to the GTM configuration. When Auto-Discovery is enabled for a BIG-IP system, the GTM system polls the LTM big3d agent at the Auto-Discovery Request Interval setting. The Auto-Discovery Request Interval setting is 30 seconds by default.

However, the GTM system does not support Auto-Discovery for BIG-IP systems and virtual servers that use network address translation; if the target BIG-IP system or any of its virtual servers employ network address translation, the GTM system disables the Auto-Discovery feature for the entire BIG-IP system. The GTM system does not attempt to automatically discover links or virtual server objects for the target server, even those virtual servers for which address translation is not configured. In addition, the GTM system does not report any messages or alerts that log the change. For example, if the

target BIG-IP virtual server IP addresses reside in a private network space, as defined by RFC1918, and are mapped to public IP addresses that are defined on a network device such as a firewall, the GTM system will silently disable the Auto-Discovery feature for the entire BIG-IP system.

If you plan to add BIG-IP systems with translated virtual servers to the GTM configuration, you must manually configure each virtual server on the GTM system. For virtual servers using address translation, you must also configure the translation address. To manually define a virtual server, perform the following procedure:

BIG-IP 10.x - 11.4.1

1. Log in to the Configuration utility.
2. Navigate to Global Traffic > Servers > Server List.
3. Click the BIG-IP system for which you will define the virtual server.
4. Click the Virtual Servers menu.
5. Click Add.
6. Type a name for the virtual server.
7. Type the IP address and port for the virtual server in the Address and Port boxes.
8. If applicable, type the translation IP address and port for the virtual server in the Translation and Translation Service Port boxes.
9. Click Create.

Objective - 2.09 Explain the necessary steps and tools to add a new LTM to a sync group

2.09 - Understand the minimal object requirements to get a sync group up

[Link to Online Topic Content](#)

Sync Group Minimum Requirements

A GTM synchronization group is a collection of multiple GTM systems that synchronize GTM configuration settings and metrics information. You must meet several minimum requirements for GTM synchronization group members to communicate and synchronize properly.

Description

For the GTM synchronization group members to properly synchronize their configuration settings, verify that the following requirements are in place:

- GTM synchronization group members must be running the same software version

A GTM device should be running the same software version as other members in the synchronization group. GTM devices that are running different software versions will not be able to communicate and properly synchronize GTM configuration and zone files.

- Synchronization parameters must be properly defined for all members

Synchronization must be enabled and each device must have the same synchronization group name.

- NTP must be configured on each device

Before you can synchronize GTM systems, you must define the Network Time Protocol (NTP) servers for all synchronization group members. Configuring NTP servers ensures that each GTM synchronization group member is referencing the same time when verifying the configuration data that needs to be synchronized.

- Port Lockdown must be set properly for the relevant self-IP addresses

Port lockdown is a security feature that specifies the protocols and services from which a self-IP address can accept traffic. F5 recommends using the Allow

Default option for self-IP addresses that are used for synchronization and other critical redundant pair intercommunications.

- TCP port 4353 must be allowed between GTM systems

GTM synchronization group members use TCP port 4353 to communicate. You must verify that port 4353 is allowed between GTM systems.

- Compatible big3d versions must be installed on synchronization group members

The big3d process runs on BIG-IP systems and collects performance information on behalf of the GTM system. For metrics collection to work properly, synchronization group members must run the same version of the big3d process.

- A valid device certificate must be installed on all members

The device certificate is used by the F5 system to identify itself to a requesting F5 client system. The default device certificate, `/config/httpd/conf/ssl.crt/server.crt`, must be installed on each sync group member.

2.09 - Explain how to add LTM to a sync group and on which host do you run bigip_add

[Link to Online Topic Content](#)

bigip_add script

The bigip_add script is an interactive script that uses the SSH protocol to exchange iQuery SSL certificates with a remote BIG-IP system. The bigip_add script appends the local GTM system's SSL certificate to the remote BIG-IP system's list of authorized certificates (contained in the /config/big3d/client.crt file). The script then appends the remote BIG-IP system's iQuery SSL certificate to the GTM system's local list of authenticated iQuery SSL certificates (/config/gtm/server.crt).

Running the bigip_add script

The bigip_add script is run from the local GTM system when adding a BIG-IP to the Wide-IP configuration.

To run the bigip_add script, log in to the command line of the GTM system and type the following command:

```
bigip_add <BIG-IP_IP_address>
```

Objective - 2.10 Explain the necessary steps and tools to add a new GTM to an existing sync group

2.10 - Describe how to add GTM to an existing deployment (add GTM to the data center, which direction to run gtm_add, how to use gtm_add)

[Link to Online Topic Content](#)

gtm_add script

The gtm_add script is an interactive script used to integrate a new GTM into a sync group that is already defined on one or more remote GTM Controllers. The script will wipe out the current configuration of the GTM on which it is run, and replace it with the same configuration of the remote GTM system in the specified sync group. The remote GTM's SSL certificates are copied to the local GTM system using the SSH protocol.

Running the gtm_add script

You can run the gtm_add script on the new GTM system that you are integrating within a network that is configured with one or more existing GTM systems. The gtm_add script will copy the remote GTM configuration to the local GTM system. The new GTM system needs to be defined in the existing GTM system's configuration prior to running the gtm_add script.

To run the gtm_add script, log in to the command line of the GTM system and type the following command:

```
gtm_add <existing_GTM_IP_address>
```

Objective - 2.11 Explain how to troubleshoot and verify sync group mesh

2.11 - Explain how to troubleshoot and verify sync group mesh

[Link to Online Topic Content](#)

If you are experiencing issues relating to GTM synchronization group communication and need to troubleshoot the issue, perform the following procedures.

Symptoms

As a result of synchronization and iQuery connection issues, you may encounter the following symptoms:

- GTM synchronization group members have configuration discrepancies.
- Server objects are marked as Unavailable (Red) or Unknown (Blue).
- Log messages report issues related to synchronization and iQuery connections.

Identifying synchronization/iQuery connections issues

GTM software includes utilities, such as the Configuration utility and tmsh utility, which you can use to identify synchronization/iQuery connection issues.

Configuration utility

The Configuration utility lists failing server objects as Offline (Red), and a failing iQuery connection as Not Connected (Red). The following table lists Configuration utility pages that display the status of GTM synchronization group members and iQuery connections:

Configuration utility page	Description	Location
Server List	Summary of defined GTM server objects	Global Traffic > Servers > Server List
Global Traffic statistics (iQuery)	Summary of iQuery statistics	Statistics > Module Statistics > Global Traffic > Statistics Type > iQuery
Global Traffic statistics (Summary)	Summary of GTM statistics	Statistics > Module Statistics > Global Traffic

tmsh utility

The tmsh utility lists failing server objects as Offline, and a failing iQuery connection as Not Connected. The following table lists tmsh utility commands that can be used to check the status of GTM synchronization group members and iQuery connections:

tmsh component	Description	Example commands
server	Summary of defined GTM server objects	tmsh show /gtm server all
iquery	Summary of iQuery statistics	tmsh show /gtm iquery all
gtm	Summary of GTM statistics	tmsh show /gtm

Verifying required configuration elements for synchronization group members

For GTM synchronization group members to communicate and synchronize properly, you must verify that certain requirements are in place. To do so, review the following check list:

Sync requirement	Description	Configuration utility location	tmsh
Software versions	Run the same software version for synchronization group members	System > Software Management	tmsh show /sys software
Sync settings	Use the same synchronization group settings for all members	System > Configuration > Global Traffic > General	tmsh list /gtm global-settings general all-properties
NTP	Configure NTP for all members	System > Configuration > Device > NTP	tmsh list /sys ntp servers
Port Lockdown	Use the Allow Default option for self-IPs that process iQuery traffic	Network > Self IPs	tmsh list /net self allow-service
iQuery port	Verify that TCP port 4353 is allowed on interconnecting devices	Not Applicable	Not Applicable
big3d versions	Run the same big3d version on all members	Not Applicable	big3d -v /shared/bin/big3d -v

Reviewing log files

Reviewing the log files is one way to determine the cause of synchronization/iQuery connection issues. Some of the logging related to synchronization/iQuery connection issues is as follows:

Device certificates messages

The BIG-IP system uses Secure Socket Layer (SSL) certificates for inter-device communication using the iQuery protocol. If device certificates are missing or expired on one of the GTM synchronization group members, the system will be marked Offline and the system logs an error message that appears similar to the following example to the `/var/log/gtm` file:

```
SSL error:14090086:SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed
```

iQuery Connectivity messages

The iQuery protocol uses TCP port 4353 to connect to synchronization group members. The system logs a successful iQuery connection to the `/var/log/gtm` file.

For example:

```
gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>
```

```
gtmd[8472]: 011ae01c:5: Connection complete to <iquery_peer>. Starting SSL handshake
```

```
gtmd[11895]: 011a5003:1: SNMP_TRAP: Server /Common/<hostname> (ip=<iquery_peer>) state change red --> green
```

```
gtmd[11895]: 011a5008:1: SNMP_TRAP: GTM /Common/<hostname> (<iquery_peer>) joined sync group default
```

If the iQuery protocol is blocked; for example, by a router ACL, or packet filter, the GTM system marks its iQuery peer as Unavailable and attempts to reestablish the iQuery connection every 10 seconds. When this behavior occurs, a log sequence appears in the `/var/log/gtm` file that appears similar to the following example:

```
gtmd[11895]: 011a500c:1: SNMP_TRAP: Box <iquery_peer> state change green --> red (Box <iquery_peer> on Unavailable)
```

```
gtmd[11895]: 011a5004:1: SNMP_TRAP: Server /Common/<hostname> (ip=<iquery_peer>) state change green --> red (No communication)
```

```
gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>
```

```
gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>
```

```
gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>
```

```
gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>
```

NTP messages

The Synchronization Time Tolerance setting specifies the number of seconds that one system clock can be out of sync with another system clock in the synchronization group. If the time difference between synchronization group members is greater than the Synchronization Time Tolerance value, the system logs a message to the `/var/log/gtm` file that appears similar to the following example:

```
gtmd[11895]: 011a0022:2: Time difference between GTM /Common/B3900-242 and me is 486 seconds --  
Make sure NTP is running and GTM times are in sync
```

This error message is an indication that NTP may not be configured on one or more synchronization group members.

Verifying big3d operation

The `big3d` process collects metrics information on behalf of the GTM system. GTM synchronization group members must run the same version of `big3d`; the running version should usually be `/shared/bin/big3d`, and the default location for the daemon is `/usr/sbin/big3d`. To verify that the version/build number is the same, perform the following procedure on all synchronization group devices:

1. Log in to the command line.
2. To verify the `big3d` version in the `/shared/bin` directory, type the following command:
`/shared/bin/big3d -v`
3. To verify the `big3d` version in the `/usr/bin` directory, type the following command:
`/usr/sbin/big3d -v`
4. If the `/shared/bin/big3d` process is older, copy the newer `big3d` process to the `/shared/bin/` directory.

For example, you would type the following command:

```
bigstart stop big3d && cp -a $(which big3d) /shared/bin/ && bigstart start big3d
```

Troubleshooting iQuery connectivity

GTM systems in a synchronization group create an iQuery mesh across synchronization group members. For example, the local GTM system's `gtmd` process opens an iQuery connection to its own `big3d` process, and to remote synchronization group member's `big3d` process. There may be occasions when you need to test iQuery connectivity between synchronization group members. For example, if log messages indicate that a GTM system has marked its iQuery peer as Unavailable, you can perform the following troubleshooting procedure to test TCP port 4353 connectivity:

1. Log in to the command line.
2. To verify the iQuery connection status, enter the following `netstat` command:

```
netstat -na |grep 4353
```

The following `netstat` output indicates that the local system (10.11.16.238) is listening on port 4353 and has an iQuery connection established to its own `big3d` process. In addition, the local system and its iQuery peer (10.11.16.242) have established an iQuery mesh:

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State
tcp	0	0	:::4353	:::*	LISTEN
tcp	0	0	::ffff:10.11.16.238:52794	::ffff:10.11.16.238:4353	ESTABLISHED
tcp	0	0	::ffff:10.11.16.238:4353	::ffff:10.11.16.242:58779	ESTABLISHED
tcp	0	0	::ffff:10.11.16.238:4353	::ffff:10.11.16.238:52794	ESTABLISHED
tcp	0	0	::ffff:10.11.16.238:46882	::ffff:10.11.16.242:4353	ESTABLISHED

3. If the synchronization group iQuery mesh is incomplete, you can use the `iqdump` command to determine if the iQuery packets arrive at the destination.

If the iQuery channel is not established, `iqdump` returns with an SSL error similar to the following example:

```
# iqdump 10.10.10.20
```

```
46947856243768:error:14090086:SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed:s3_clnt.c:1168:
```

If the iQuery channel is established, `iqdump` returns XML similar to the following example:

```
# iqdump 10.10.10.20
```

```
<!-- Local hostname: lc1.example.com -->
```

```
<!-- Connected to big3d at: ::ffff:10.10.10.10:4353 -->
```

```
<!-- Subscribing to syncgroup: default -->
```

```
<!-- Tue May 6 09:55:43 2014 -->  
  
<xml_connection>  
  
<version>11.5.1</version>  
  
<big3d>big3d Version 11.5.1.0.0.110</big3d>
```

Verifying device SSL certificates

Each synchronizing group member must have a valid SSL device certificate installed in the `/config/httpd/conf/ssl.crt/` directory for iQuery connections to succeed. If log messages indicate an issue with a device certificate on one of the synchronization group members, you can verify the certificate status by performing the following procedure:

1. Log in to the command line.
2. Check the status of the device certificate by typing the following command:
`openssl x509 -noout -text -in /config/httpd/conf/ssl.crt/server.crt`
3. Verify the certificate validity date and confirm whether the certificate is expired.
4. If necessary, renew the certificate. To do so, refer to SOL6353: Updating an SSL device certificate on a BIG-IP system.

Troubleshooting daemons

The `tmm`, `mcpd`, `big3d`, and `gtmd` processes are all critical to synchronizing GTM configurations. To confirm that the daemons are running as expected, use the `bigstart` command.

For example, to confirm the status of the `tmm`, `mcpd`, `big3d`, and `gtmd` processes, type the following command:

```
bigstart status tmm mcpd big3d gtmd
```

If the `mcpd` process is consuming more than 90 percent of a CPU, and synchronizing actions such as saving the configuration may fail. To check the CPU utilization for the `mcpd` process, type the following command:

```
top -p `pidof mcpd`
```

To quit, type q.

Objective - 2.12 Explain the use of device certificates in iQuery [SSL components, expiration, 3rd party certs]

2.12 - Explain the implications of device certificate expiration

[Link to Online Topic Content](#)

Certificates

The BIG-IP system uses SSL encryption for securing administrative connections. Each synchronizing group member must have a valid SSL device certificate installed in the `/config/httpd/conf/ssl.crt/` directory for iQuery connections to succeed. Third party certificates can be used for this function. To determine the expiration date for SSL certificate and key pairs, you can use either the `tmsh` utility or the BIG-IP Configuration utility.

If the BIG-IP system is communicating with a GTM system, you must ensure that the renewed device certificate is added to the trusted device certificate file on the GTM system. You can do so by running the `bigip_add` utility on the GTM system. Failure to run the `bigip_add` utility prevents the GTM system from communicating with this BIG-IP system.

Note: If the GTM systems are in a synchronization group, you need to perform this procedure on only one GTM system. If the GTM systems are not in a synchronization group, you must perform this procedure on each GTM system that communicates with this BIG-IP system that had its device certificate updated or changed.

This procedure authenticates, copies, and adds the certificate to the `/config/gtm/server.crt` file, which allows the GTM system to authenticate the certificate each time it communicates with this BIG-IP system. To add the renewed device certificate, perform the following procedure on the GTM system:

Impact of procedure: None.

1. Log in to the GTM command line.
2. Enter the following command: **bigip_add <BIG-IP_IP_address>**
3. When prompted, enter the root password of the BIG-IP system.

Objective - 2.13 Explain how to verify listener responses

2.13 - Including static versus intelligent, TTLs, number of answer records, stats in the profile, logging, and tcpdump

[Link to Online Topic Content](#)

Listener Responses

Static vs Intelligent - A static resolution is when the listener is simply returning a A record or AAAA response from a Zone file whether it is Bind on the GTM or bridged to a DNS server in the network. (Some may consider the fallback mode as a static resolution but a fallback only occurs after the other resources for the Wide-IP have failed) An intelligent response would be a response to a Wide-IP request that is deciding best IP for resolution based on the status of the Wide-IP resources.

[Link to Online Topic Content](#)

Time-to-Live (TTL) – Each resource in GTM has an associated time-to-live (TTL) value. A TTL is the amount of time (measured in seconds) for which the system considers metrics valid. The timer values determine how often Global Traffic Manager refreshes the information.

Each resource also has a timer value. A timer value defines the frequency (measured in seconds) at which Global Traffic Manager refreshes the metrics information it collects. In most cases, the default values for the TTL and timer parameters are adequate. However, if you make changes to any TTL or timer values, keep in mind that an objects TTL value must be greater than its timer value.

Number of Answer Records – I am not sure what they are looking for here other than that if you are querying for a Wide-IP name that is using an algorithm like round robin you will see multiple records returned according to the algorithm not just a static IP address every time. Thus you could see that the Wide-IP is working correctly.

[Link to Online Topic Content](#)

Statistics - You can access Global Traffic Manager statistics in two ways:

- Through the Statistics option on the Main tab of the navigation pane
- Through the Statistics menu from various main screens for different components

Both methods take you to the same screen within Global Traffic Manager. When you access statistics through a menu on the main screen for a given network component,

the Statistics screen is pre-configured for the given network element, although you can switch to a different set of statistics at any time.

Additionally, you can use the search feature to locate a specific component or group of components. The default search value is an asterisk (*), which instructs the system to display all relevant components in a list. You can type a string in the box, and when you click the Search button, the system modifies the list to show only those components that match the string.

Tip: You can also access statistics from the command line using the tmsh command show.

[Link to Online Topic Content](#)

Logging – When GTM receives a DNS name resolution request for a Wide-IP, in order to send a response; the system makes a load-balancing decision. The decision is based on the load-balancing method configured on the Wide-IP, the number of pools associated with the Wide-IP, and the applicable number of members in each pool.

You can see information about how GTM made the load-balancing decision in the logs; reviewing the logs can help determine how to fine-tune your network. When you want to view the global server load-balancing decisions made by GTM in the high-speed remote logs, configure the verbosity of the information that displays in the logs.

Tcpdump – The tcpdump tool is available on all BIG-IP platforms. As a packet capture tool you can see all requests and responses from the BIG-IP platform. Capturing traffic from the Listener IP address is one way to verify responses.

Objective - 2.14 Explain how to verify that DNSSEC is working

2.14 - Including records getting signed, authoritative bit set, sig files in correct location

[Link to Online Topic Content](#)

Content

We use a test client to access the GTM Wide-IP to perform DNS lookup requests. A DNS client application called Dig can be used to query the DNS Server.

Launch a terminal application and issue a request that includes DNSSEC, such as:

```
dig @bigip10.siterequest.com +dnssec +multiline www.dnssec.f5demo.com
```

You see a result similar to the following example:

```
; <<> DiG 9.6.0-APPLE-P2 <<> @bigip10.siterequest.com +dnssec +multiline www.dnssec.f5demo.com
; (1 server found) ;; global options: +cmd ;; Got answer: ;; ->>HEADER<<- opcode: QUERY, status:
NOERROR, id: 60496 ;; flags: qr aa rd ad; QUERY: 1, ANSWER: 6, AUTHORITY: 0, ADDITIONAL: 1 ;;
WARNING: recursion requested but not available
```

```
;; OPT PSEUDOSECTION: ; EDNS: version: 0, flags: do; udp: 4096 ;; QUESTION SECTION:
;www.dnssec.f5demo.com. IN A
```

```
;; ANSWER SECTION:
```

```
www.dnssec.f5demo.com.          30 IN A 65.197.145.93
```

```
www.dnssec.f5demo.com.          30 IN RRSIG A 7 4 30 20100116005323 (
20100109005323 31052 dnssec.f5demo.com.
NtOnSwwK1JhbYgsCY5EhVSzZ7475A6NAfcaAnhXkiYCN
us+0TYKoRwXfGKodNjd/wjrcD+J08Vz8SxSuQ19cY9Jx
Kt01o7ghLgvcIemyYTsICEWXJ98FrX9MdJcQvaeg3Qvj
FKQMVHvrNxVgzTkTdcVvK8Q/zgVMCbejcEK29iI= )
```

```
www.dnssec.f5demo.com.          30 IN RRSIG A 7 4 30 20100116005323 (
20100109005323 61232 dnssec.f5demo.com.
vJS+4Cf8EM6b73LG6Lb1xxNxENWx7y1ct7QdggCnCS1u
9iD0pW0dDKaZIH8ya4UD8Ar/V+yJjrPxA25Shk/nh1W4t
81/R+njx1MJJoZ9a71Y8cHMqXLPyGepYXVHY70J+akp83
3oYbFbMvg7YbnYEItNUEM+6LuitXo89FUTaY2QI= )
```

```
www.dnssec.f5demo.com.          30 IN RRSIG A 7 4 30 20100116005323 (
20100109005323 46472 dnssec.f5demo.com.
fdio5eNraa1eBM+/NCbVT6rKWukoq1Z2VICpY2wa2X/Q
ocWRcyO1da2s1pKEh6LRTEZ4z13MrwQbyh6AuaaU/LEZ
8VEU2ViK90wwKBLMFsnWqPMYLZ0PSd3a+ANcbr869vsJ
9F4DSs9CfbVJdOkaGFqPYWjwpqMLxN/B1aH1Npw= )
```

```
www.dnssec.f5demo.com.          30 IN RRSIG A 7 4 30 20100116005323 (
20100109005323 64235 dnssec.f5demo.com.
7cpHDxhdqAips+rLTpprDnj5Jc+J6qDZ6x9JNYR4Pe1J
```

```
Mp1pmVq72tYUVIcJPZ3fpdpCW83cLSj6Ij83/zPORP3p
MubfIe4mtk3ysGQGzA/Aatx8+J3T8AHHi00y7qo4XEUY
N1sItDAi9nCX1XD4QwBXmQtur+QYESQCy937uRM= )
```

www.dnssec.f5demo.com.

```
30 IN RRSIG A 7 4 30 20100116005323 (
20100109005323 28328 dnssec.f5demo.com.
K2WXvNNMa4AEGE8q5e7qPcdg9ki0LcMgOgiHhwG8fD5K
qfLaqo89BNdhba12AKs+F/8T+H0K5ZNRnW/L591vTFxT
A15iVEzZwO9Uv008UeZtvWafYbfq41D6e/S0KjnXo2kR
W3DiNSA2UFC1QSNp5Aic+cF0IKEem/yJ/+PwxmQ= )
```

```
;; Query time: 70 msec ;; SERVER: 65.197.145.83#53(65.197.145.83) ;; WHEN: Fri Jan 8 16:53:23
2010 ;; MSG SIZE rcvd: 1077
```

Objective - 2.15 Given a scenario explain how to validate system health for proper operation

2.15 - Given a scenario explain how to validate system health for proper operation

[Link to Online Topic Content](#)

Validate System Health

To see additional platform performance information, use the following steps:

In version 11.x of the BIG-IP Configuration Utility:

1. Click Statistics.
2. Click Performance.

All categories are shown under the All tab or you can see the break outs of System, Connections, Throughput and Cache.

Each category will provide MRTG based graphs that display information about how the Global Traffic Manager is performing. You can use this information to help you determine how to modify the configuration to obtain the best possible performance from the system.

Viewing performance data

The Global Traffic Manager provides two types of performance data graphs on the performance screen: the GTM Performance and GTM Request Breakdown graphs. You can view detailed versions of each graph by clicking the View Detailed Graph link.

About the GTM Performance graph

The GTM Performance graph shows the throughput of the Global Traffic Manager. The graph includes the following data:

GTM Requests

Represents the number of incoming DNS requests.

GTM Resolutions

Represents the number of incoming DNS requests that were resolved by any method.

GTM Resolutions Persisted

Represents the number of incoming DNS requests that were resolved by a persistence record.

GTM Resolutions Returned to DNS

Represents the number of incoming DNS requests that were not resolved by the Global Traffic Manager, but were instead passed on to the DNS server for resolution.

About the GTM Request Breakdown graph

The GTM Request Breakdown graph includes the following data:

GTM Type A - IPv4 Requests

Represents IPv4-formatted requests.

GTM Type AAAA/A6 - IPv6 Requests

Represents IPv6-formatted requests.

To view performance data

1. On the Main tab of the navigation pane, expand Overview and then click Performance.

The Performance screen opens.
2. On the menu bar, click Global Traffic.

The Performance screen displays the global traffic management Graphs.
3. From the Graph Interval list, select the time period for which you want to view performance data.
4. Click the Refresh button to update the graphs.
5. Click the View Detailed Graph links to view the detailed graphs.

Section 3 – Operations and Troubleshooting

Objective - 3.01 Given a scenario determine the impact of software updates in a group on monitoring and configuration state

3.01 - Given a scenario determine the impact of software updates in a group on monitoring and configuration state

[Link to Online Topic Content](#)

Scenario Based Questions

This topic is focused on the installation of software updates or hotfixes to the GTM and the impact to production traffic. Experience with applying software updates to a GTM in environment will give the candidate the ability to answer the questions on this topic.

To prepare for scenario based questions the candidate will need to complete hands-on configuration and testing of the configuration on the GTM. This will allow the candidate to better understand how different configurations can produce different results. All F5 exams use scenario-based questions that make the candidate apply what they know to a situation to determine the resulting outcome.

During the upgrade, the unit you are performing the task on will be unavailable to resolve any DNS requests. In addition, no configuration changes should be introduced to the GTM synchronization group until all members of the synchronization group have been upgraded to the same version.

If the health monitoring that is configured in the Sync group is currently using the GTM, that is being upgraded, to do some of the monitoring directly, then that GTM device will not be able to do monitoring during the upgrade. Another device in the sync group will perform the monitoring task while the system is off line. If there is a firewall between the other GTM devices and the monitored devices there must be policy that allows the monitoring to take place. Otherwise monitoring of the devices will fail and there could be an impact to production traffic.

F5 recommends that you perform the following steps to ensure a successful GTM software version upgrade and/or hotfix installation:

GTM pre-upgrade recommendations

- Before upgrading a single GTM system or a group of GTM systems, F5 recommends that you perform all of the following steps before you apply the upgrade. These actions ensure that the systems are ready to receive the upgrade and minimize downtime while the upgrade is being installed.
- Verify that all trusted device certificates and local certificates that are used for secure communication between GTM devices are up to date and will not expire during the maintenance period, or shortly following the maintenance period.
- Create a current backup copy of the local configuration for each GTM system, and store the backup in a secure and separate location from the device.
- Download the installation IM and/or hotfix to each GTM device and verify that the MD5 checksum of the hotfix file is correct.



Objective - 3.02 Given a scenario determine what is the effect of changing the features enabled in a DNS profile

3.02 - Including enabling/disabling recursion, protocol, unhandled query behavior, and making sure BIND is not enabled in the profile or in the GTM pools, etc.

[Link to Online Topic Content](#)

Recursion

When a DNS server sets the recursion available (RA) bit in a DNS response, the DNS server is indicating to the client that it will query other name servers for requested domain names if the DNS server's zone files do not contain the answer. By default, DNS recursion is disabled on GTM systems. Under certain circumstances, you may want to enable DNS recursion on the GTM system.

After DNS recursion is enabled on the GTM system, you can create an ACL to limit which IP addresses or network addresses are allowed to make recursive queries to the GTM system.

Enabling DNS recursion on the GTM system

1. Log in to the Configuration utility.
2. Navigate to DNS > Zones > ZoneRunner > named Configuration.
3. In the named Options window, locate the options section of the named.conf file and change the recursion statement to the following:
recursion yes;
4. Click Update.

[Link to Online Topic Content](#)

Profile settings

General property	Description	Default Value
Name	Specifies the user-supplied name of the profile. Specifying a name for your profile is required.	No default value
Parent Profile	Specifies the profile from which your custom profile is derived.	dns
Global Traffic Management	Specifies whether the system uses Global Traffic Manager to manage the response.	Enabled
DNS IPv6 to IPv4	<p>Specifies whether you want the BIG-IP system to convert IPv6-formatted IP addresses to IPv4-formatted IP addresses. The possible values are: Disabled: Indicates that the BIG-IP system does not map IPv4 addresses to IPv6 addresses.</p> <p>Secondary: Indicates that the BIG-IP system receives an AAAA query and forwards the query to a DNS server.</p> <p>Immediate: Indicates that the BIG-IP system receives an AAAA query and forwards the query to a DNS server.</p> <p>v4 Only: Indicates that the BIG-IP system receives an AAAA query, but forwards an A query to a DNS server.</p> <p>When you select Secondary, Immediate, or v4 Only, you must also provide a prefix in the IPv6 to IPv4 Prefix field and make a selection from the IPv6 to IPv4 Additional Section Rewrite list.</p>	Disabled
DNS Express	Indicates whether the dns-express service is enabled. The service handles zone transfers from the primary DNS server.	Enabled
DNSSEC	Specifies whether the system signs responses and replies to DNSSEC-specific queries (for example, DNSKEY query type).	Enabled

Unhandled Query Actions	<p>Specifies whether the system uses the local BIND server on the BIG-IP system. The possible values are:</p> <p>Allow: Indicates that the BIG-IP system forwards the connection request to another DNS server or DNS server pool. If a DNS server pool is not associated with a listener and the Use BIND Server on BIG-IP setting is set to Enabled, connection requests are forwarded to the local BIND server.</p> <p>Drop: Indicates that the BIG-IP system does not respond to the query.</p> <p>Reject: Indicates that the BIG-IP system returns the query with the REFUSED return code.</p> <p>Hint: Indicates that the BIG-IP system returns the query with a list of root name servers.</p> <p>No Error: Indicates that the BIG-IP system returns the query with the NOERROR return code.</p>	Allow
Use BIND Server on BIG-IP	<p>Specifies whether the system forwards non-Wide-IP queries to the local BIND server on the BIG-IP system. For best performance, disable this setting when using a DNS cache.</p>	Enabled
Process Recursion Desired	<p>Indicates whether to process client-side DNS packets with Recursion Desired set in the header. If set to Disabled, processing of the packet is subject to the unhandled-query-action option.</p>	Enabled
DNS Cache	<p>Specifies whether the system caches DNS responses.</p> <p>Enabled: Indicates the BIG-IP system caches DNS responses handled by the virtual servers associated with this profile. When you enable this setting, you must also select a cache from the DNS Cache Name list.</p> <p>Disabled: Indicates the BIG-IP system does not cache DNS responses handled by the virtual servers associated with this profile. However, the profile retains the association with the DNS cache in the DNS Cache Name field. Disable this setting when you want to debug the system.</p>	Disabled
DNS Cache Name	<p>Specifies the user-created cache that the system uses to cache DNS responses. When you select a cache for the system to use, you must also enable the DNS Cache setting.</p>	No default value
DNS Security	<p>Indicates whether DNS firewall capability is enabled.</p>	Disabled
DNS Security Profile Name	<p>Specifies the DNS security profile to use.</p>	No default value

Logging	Indicates whether to enable high speed logging for DNS queries and responses. When it is set to Enabled, you must also specify a Logging Profile.	Disabled
Logging Profile	Specifies the DNS logging profile used to configure what events get logged and their message format. These are the DNS Logging profiles you create using the Other option under Profiles in the BIG-IP Configuration utility.	No default value

Objective - 3.03 Explain how to renew device certificates and update them in the sync group

3.03 - Explain how to renew device certificates and update them in the sync group

[Link to Online Topic Content](#)

Renewing the self-signed device certificate on the BIG-IP system

The BIG-IP system uses the device certificate to authenticate access to the Configuration utility, and to accommodate device-to-device communication processes, such as ConfigSync, big3d, and gtmd.

Note: Starting in BIG-IP 10.2.4 and 11.2.0, the device certificate that ships with the BIG-IP system is valid for 10 years.

To renew the self-signed device certificate on the BIG-IP system, perform the following procedures:

Configuration utility

Note: Each F5 device has a unique x509 device certificate that you need to renew independently. This policy applies to standalone systems and redundant pair members.

Impact of procedure: Renewing the device certificate will require BIG-IP Configuration utility users to re-authenticate.

1. Log in to the BIG-IP Configuration utility.
2. Navigate to System > Device Certificates.
3. Click Renew.
4. From the Issuer box, select Self.
5. From the Country box, select the appropriate country.

Note: F5 recommends that you configure a unique Common Name for each GTM synchronization group member. If this is not possible, you must have at least one other field (typically Division) with a unique value for each GTM system in a synchronization group to prevent synchronization issues.

6. Click Finished.
7. (GTM and Link Controller only) After renewing the BIG-IP device certificate on a GTM or BIG-IP Link Controller system, you must ensure that the system copies the new certificates to the trusted device certificate file on remote F5 devices

with which the GTM or Link Controller system communicates. You can exchange the device certificates by running the `bigip_add` utility from the command line of each GTM or Link Controller system containing a new certificate.

To exchange device certificates with all the F5 devices listed in the `/config/bigip_gtm.conf` file (`/config/gtm/wideip.conf` in BIG-IP 9.x through 10.x), run the following command: `bigip_add`



Objective - 3.04 Explain the impact of restoring a UCS on a GTM

3.04 - Including how to restore a GTM after an RMA and the effect on zone files

[Link to Online Topic Content](#)

UCS Restore on a GTM

A user configuration set (UCS) is a backup file that contains BIG-IP configuration data that can be used to fully restore a BIG-IP system in the event of a failure or Return Materials Authorization (RMA) replacement.

The UCS archive, by default, contains all of the files that are required to restore your current configuration to a new system, including configuration files, the product license, local user accounts, and Secure Socket Layer (SSL) certificate/key pairs.

Before installing a UCS archive on a GTM system that will be added to an existing sync group, note the following information:

- When the GTM system loads the UCS, the MCP daemon generates a new configuration change identifier (commit ID).
- The new commit ID causes the GTM system to synchronize the contents of the UCS archive to the GTM sync group.

You may like this behavior if you want to install the UCS on a current sync group member and roll the entire GTM sync group back to the configuration contained in the UCS archive.

However, in some cases you may want to install a UCS and prevent the GTM system from synchronizing the contents of the UCS archive to the GTM sync group. For example, if you are installing a UCS on a new device, such as an RMA replacement, you can prevent the system from synchronizing the configuration in the UCS to the sync group.

To prevent synchronization when installing a UCS archive on a GTM system use the following procedure:

Impact of procedure: The following procedure requires that you temporarily disconnect the GTM system's TMM switch port interfaces from the production network and leave the management interface connected. The goal is to isolate the GTM system from the production network and restore the UCS file. This action prevents the system from potentially synchronizing an older configuration to the sync group.

1. Physically disconnect the GTM system's TMM switch port interfaces from the network.
2. Log in to the Configuration utility using the management interface.
3. Restore the UCS archive.

Note: For more information, refer to SOL13132: Backing up and restoring BIG-IP configuration files (11.x).

4. After the UCS archive is installed, navigate to System > Configuration > Global Traffic > General.
5. Clear the Synchronization check box.
6. Click Update.
7. Reconnect the GTM system's TMM switch port interfaces to the network.
8. Log in to the GTM command line.
9. Add the GTM system to the GTM synchronization group by typing the following command:

```
gtm_add <IP address of a member of the target GTM synchronization group>
```

Note: The gtm_add script is interactive, and replaces the GTM configuration on the system on which it is run with the configuration of the remote GTM system in the specified sync group.

Objective - 3.05 Explain the importance of running compatible versions of big3d on the LTM and GTM

3.05 - Explain how to update big3d on LTM (big3d_install) and what concerns might be when EM is also updating GTM

[Link to Online Topic Content](#)

big3d

The big3d process runs on all BIG-IP systems, and provides metrics collection data for BIG-IP systems. The big3d_install script is an interactive script that allows you to install the current version of the big3d process on remote F5 systems. If the current or newer version of the big3d process is found to be running on the remote BIG-IP system, installation is skipped for that BIG-IP system. The big3d_install script also copies the trusted device certificate from the local GTM system to the /config/big3d/client.crt file on the remote BIG-IP system, and the trusted server certificate from the remote BIG-IP system to the /config/gtm/server.crt file on the local GTM system.

big3d version management

To facilitate proper iQuery communication in your environment, you should be aware of the following big3d version management information:

- F5 recommends that all devices communicating over iQuery run the same big3d version
- GTM synchronization group communication

GTM sync group members are required to run the same big3d version.

- GTM/BIG-IP communication

Monitored BIG-IP systems must run the same or newer big3d version as the GTM devices that are monitoring them.

- Enterprise Manager/BIG-IP communication

Managed BIG-IP systems must run the same or newer big3d version as the Enterprise Manager devices that are collecting data from them. Upgrading or downgrading big3d from the Enterprise Manager to a different version from the GTM will cause the GTM to temporarily change the status to Offline for all objects hosted on the system being upgraded.



big3d installation behavior

The big3d process initiates and runs as follows on BIG-IP and Enterprise Manager devices:

- The default big3d process is located in the /usr/sbin directory.
- At start up, the BIG-IP system verifies that the /shared/bin/big3d process exists.

If the /shared/bin/big3d process exists and has an older modification time than the /usr/sbin/big3d process, the system copies the /usr/sbin/big3d process to the /shared/bin directory and runs the /shared/bin/big3d process.

If the /shared/bin/big3d process does not exist, the system copies the /usr/sbin/big3d process to the /shared/bin directory and runs the /shared/bin/big3d process.

The result should be that the BIG-IP system runs the newer version of the big3d process from the /shared/bin directory. However, under certain conditions, the big3d process in the /shared/bin directory may be a different version than the instance in the /usr/sbin directory. For example, this may occur if an Enterprise Manager device pushes an older version of the big3d process to the system.

big3d_install

The big3d_install script attempts to use an iQuery connection over TCP port 4353 to copy the certificates and big3d process to the remote BIG-IP systems. The script uses SSH if the iQuery connection fails.

Running the big3d_install script

To run the big3d_install utility, log in to the command line of the GTM system and type the following command:

```
big3d_install <BIG-IP_IP_address>
```

Note: If no IP addresses are specified, the script will attempt to install the current version of the big3d process on all the BIG-IP controllers listed in the bigip_gtm.conf file.

Objective - 3.06 Explain how to properly add/remove device from iQuery mesh

3.06 - Explain how to properly add/remove device from iQuery mesh

[Link to Online Topic Content](#)

Content

A GTM synchronization group synchronizes the GTM configuration and metrics among its members.

If you attempt to remove a member from the GTM synchronization group by changing the name of the GTM synchronization group for that member, the new name will be synchronized to the remaining members instead.

To properly remove a member from the GTM synchronization group, clear the Synchronization and Synchronize DNS Zone Files check boxes in the Configuration utility, or set the synchronization and synchronize-zone-files options to no in the tmsh utility.

Removing a GTM system using the Configuration utility (10.x - 11.4.1)

1. Log in to the Configuration utility of the GTM system that you want to add to the sync group.
2. Navigate to System > Configuration > Global Traffic > General.
3. Clear the Synchronization check box and the Synchronize DNS Zone Files check box.
4. Click Update.

To re-add the member to its previous GTM synchronization group, use the gtm_add utility.

Objective - 3.07 Explain the effect of adding a resource record without using ZoneRunner

3.07 - Explain the effect of adding a resource record without using ZoneRunner

[Link to Online Topic Content](#)

Directly Editing a Resource Record In Bind

The GTM ZoneRunner utility uses dynamic update to make zone changes. All changes made to a zone using dynamic update are written to the zone's journal file. When the GTM system restarts after a shutdown, the system replays the journal file to incorporate any updates that took place after the last zone file update into the zone. Dynamic update periodically flushes the complete contents of the updated zone to its zone file and automatically deletes the journal file. However, if manual updates to a zone are required, the zone files must be frozen to prevent dynamic updates from occurring and overwriting changes to the zone file.

Important: F5 recommends that the ZoneRunner utility manages the DNS/BIND file, rather than manually editing the file. If you are required to manually edit the zone files, you must freeze the zone files to avoid issues with name resolution and dynamic updates.

ZoneRunner can be stopped and an individual zone or all zones can be frozen. While the zones are frozen, dynamic updates cannot occur, but normal name resolution is allowed.

Important: To prevent the journal files from being synchronized if the GTM is configured to synchronize DNS zone files, the zone must be frozen on all GTM systems.

3.07 - Explain how to maintain zones via ZoneRunner, including moves, adds, and deletions

[Link to Online Topic Content](#)

Using ZoneRunner

One of the modes in which you operate the GTM is the node mode. In node mode, the GTM is responsible not only for load balancing name resolution requests and monitoring the health of your physical and logical network; it is also responsible for maintaining the DNS zone files that map name resolution requests to the appropriate network resource.

The ZoneRunner utility is available in GTM. The ZoneRunner utility is a zone file management utility that can manage both DNS zone files and your BIND configuration. With the ZoneRunner utility, you can create, modify, and delete zone files and manage the records in the zone files. Additionally, you can transfer zone files to another name server, or import zone files from another name server.

Note: To see the procedures to follow to do all of these tasks please see the link at the beginning of this section.

Objective - 3.08 Explain the effects and implications of securing/hardening with respect to normal operation, iQuery and resolution

3.08 - Including port lockdown, packet filters, iQuery, SSH, effects of appliance mode on LTM, bridge GTM, and the limitations of not having advanced shell access to GTM

[Link to Online Topic Content](#)

Content

Securing your BIG-IP platform is a critical part of configuration. Remembering that a complete port lockdown will typically break iQuery communications between the different BIG-IP platforms, as well as break DNS bridge mode if configured. Every scenario will be a little bit different in how you allow/restrict communications to the BIG-IP platform. And even though you may have setup the BIG-IP correctly in respect to your architectural needs, other restrictive devices on the network like firewalls can block your devices from working correctly.

[Link to Online Topic Content](#)

Network Access Management

The Port Lockdown feature controls network ports that are accessible on a self-IP. By default, the BIG-IP system allows access to only a limited set of the available ports, and the default set includes those ports required for administrative access and inter-device communication, such as in a high-availability configuration.

For example:

- Port 4353 is used in GTM deployments to transfer sync-group data.
- Port 443 allows the Configuration utility to be accessed on a Traffic Management Microkernel (TMM) switch interface, in addition to the MGMT interface.

The default Port Lockdown configuration and information on modifying the Port Lockdown settings for a given interface is available in SOL7317: Overview of port lockdown behavior.

While some ports may need to be open to ensure a properly-functioning configuration, access to these ports can be further controlled by the use of packet filters. Packet filters

allow you to control access based on combinations of criteria that include source IP address, destination IP address, MAC address, and so on.

[Link to Online Topic Content](#)

Packet Filtering

Packet filters enhance network security by specifying whether a BIG-IP system interface should accept or reject certain packets based on criteria that you specify. Packet filters enforce an access policy on incoming traffic. They apply to incoming traffic only.

Packet filter functions can be replaced by an upstream firewall but that means you are releasing administrative control of the restrictions that can be applied by packet filters.

[Link to Online Topic Content](#)

Appliance Mode

Beginning with the release of BIG-IP 10.2.1-HF3, BIG-IP systems now have the option of running in Appliance mode. Appliance mode is designed to meet the needs of customers in especially sensitive sectors by limiting the BIG-IP system administrative access to match that of a typical network appliance and not a multi-user UNIX device.

Technical Restrictions in Appliance mode

- Access to the bash shell has been removed.
- Administrative access is limited to the Configuration utility, bigpipe shell (bpsh), and the Traffic Management Shell (tmsh).
- The root user cannot log in to the device by any means, including the serial console.
- On platforms that include the Always-On Management (AOM) subsystem, the AOM is not able to access the host. The AOM is only able to reset the host using a hardware reset command.
- On VIPRION platforms, SSH access between blades is not allowed by way of the ssh slot<X> command syntax. You can SSH to only the management IP address of each blade.
- Once you have enabled Appliance mode, you cannot disable it in any way other than obtaining a new license from F5 and performing a clean installation of the software.

Objective - 3.09 Identify GTM specific command line tools and TMSH GTM specific commands

3.09 - Identify GTM specific command line tools and TMSH GTM specific commands

[Link to Online Topic Content](#)

GTM specific command line tools

Reviewing the log files is one way to determine the cause of synchronization/iQuery connection issues.

GTM logs can be viewed at command line by typing the following command:

```
more /var/log/gtm
```

Device certificates messages

The BIG-IP system uses Secure Socket Layer (SSL) certificates for inter-device communication using the iQuery protocol. If device certificates are missing or expired on one of the GTM synchronization group members, the system will be marked Offline and the system logs an error message that appears similar to the following example to the `/var/log/gtm` file:

```
SSL error:14090086:SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed
```

iQuery Connectivity messages

The iQuery protocol uses TCP port 4353 to connect to synchronization group members. The system logs a successful iQuery connection to the `/var/log/gtm` file. For example:

```
gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>
gtmd[8472]: 011ae01c:5: Connection complete to <iquery_peer>. Starting SSL handshake
gtmd[11895]: 011a5003:1: SNMP_TRAP: Server /Common/<hostname> (ip=<iquery_peer>) state change red --> green
gtmd[11895]: 011a5008:1: SNMP_TRAP: GTM /Common/<hostname> (<iquery_peer>) joined sync group default
```

If the iQuery protocol is blocked; for example, by a router ACL, or packet filter, the GTM system marks its iQuery peer as Unavailable and attempts to reestablish the iQuery connection every 10 seconds. When this behavior occurs, a log sequence appears in the `/var/log/gtm` file that appears similar to the following example:

```

gtmd[11895]: 011a500c:1: SNMP_TRAP: Box <iquery_peer> state change green --> red
(Box <iquery_peer> on Unavailable)

gtmd[11895]: 011a5004:1: SNMP_TRAP: Server /Common/<hostname> (ip=<iquery_peer>) state
change green --> red (No communication)

gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>

gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>

gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>

gtmd[8472]: 011ae020:5: Connection in progress to <iquery_peer>

```

NTP messages

The Synchronization Time Tolerance setting specifies the number of seconds that one system clock can be out of sync with another system clock in the synchronization group. If the time difference between synchronization group members is greater than the Synchronization Time Tolerance value, the system logs a message to the `/var/log/gtm` file that appears similar to the following example:

```

gtmd[11895]: 011a0022:2: Time difference between GTM /Common/B3900-242 and me is 486
seconds -- Make sure NTP is running and GTM times are in sync

```

tmsh commands to show components

tmsh component	Description	Example commands
server	Summary of defined GTM server objects	<code>tmsh show /gtm server all</code>
iquery	Summary of iQuery statistics	<code>tmsh show /gtm iquery all</code>
gtm	Summary of GTM statistics	<code>tmsh show /gtm</code>

tmsh commands to show sync requirements

Sync requirement	Description	tmsh
Software versions	Run the same software version for synchronization group members	<code>tmsh show /sys software</code>
Sync settings	Use the same synchronization group settings for all members	<code>tmsh list /gtm global-settings general all-properties</code>
NTP	Configure NTP for all members	<code>tmsh list /sys ntp servers</code>
big3d versions	Run the same big3d version on all members	<code>/shared/bin/big3d -v</code>

3.09 - Show a GTM iQuery

[Link to Online Topic Content](#)

Troubleshooting iQuery

GTM systems in a synchronization group create an iQuery mesh across synchronization group members. For example, the local GTM system's `gtmd` process opens an iQuery connection to its own `big3d` process, and to remote synchronization group member's `big3d` process. There may be occasions when you need to test iQuery connectivity between synchronization group members. For example, if log messages indicate that a GTM system has marked its iQuery peer as Unavailable, you can perform the following troubleshooting procedure to test TCP port 4353 connectivity:

1. Log in to the command line.
2. To verify the iQuery connection status, enter the following netstat command:

```
netstat -na |grep 4353
```

The following netstat output indicates that the local system (10.11.16.238) is listening on port 4353 and has an iQuery connection established to its own `big3d` process. In addition, the local system and its iQuery peer (10.11.16.242) have established an iQuery mesh:

Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State
tcp	0	0	:::4353	:::*	LISTEN
tcp	0	0	::ffff:10.11.16.238:52794	::ffff:10.11.16.238:4353	ESTABLISHED
tcp	0	0	::ffff:10.11.16.238:4353	::ffff:10.11.16.242:58779	ESTABLISHED
tcp	0	0	::ffff:10.11.16.238:4353	::ffff:10.11.16.238:52794	ESTABLISHED
tcp	0	0	::ffff:10.11.16.238:46882	::ffff:10.11.16.242:4353	ESTABLISHED

3. If the synchronization group iQuery mesh is incomplete, you can use the `iqdump` command to determine if the iQuery packets arrive at the destination.

If the iQuery channel is not established, `iqdump` returns with an SSL error similar to the following example:

```
# iqdump 10.10.10.20
```

```
46947856243768:error:14090086:SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed:s3_clnt.c:1168:
```

If the iQuery channel is established, `iqdump` returns XML similar to the following example:

```
# iqdump 10.10.10.20
```

```
<!-- Local hostname: lc1.example.com -->  
<!-- Connected to big3d at: ::ffff:10.10.10.10:4353 -->  
<!-- Subscribing to syncgroup: default -->  
<!-- Tue May 6 09:55:43 2014 -->  
  
<xml_connection>  
  
<version>11.5.1</version>  
  
<big3d>big3d Version 11.5.1.0.0.110</big3d>
```



Objective - 3.10 Given a scenario determine what information needs to be provided when making a support call

3.10 - Given a scenario determine what information needs to be provided when making a support call

[Link to Online Topic Content](#)

Support Information

The following content describes what information is typically gathered when opening a support call. Scenario based questions usually test your knowledge by making you apply what you know to a situation. You should get familiar with the type of data that is gathered in the files below and also be familiar with using iHealth.

Provide the following information when you open a case with F5 Technical Support:

1. A full description of the issue, including the following details:
 - The symptoms of the issue, including a brief description of any systems applicable to the configuration
 - The approximate time the issue first occurred
 - The number of times the issue has recurred
 - Any error output provided by the system
 - Steps to reproduce the issue
 - Any changes you made to the system before the issue first occurred
 - Any steps you took to attempt to resolve the issue
 - Whether this is a new implementation
 - GTM:

How many datacenters and devices are applicable to the configuration?

Which devices does the issue affect?

- Whether you have already uploaded a qkview to the iHealth portal
 - (Ensure that any qkviews uploaded to the iHealth portal are linked to the support case)
2. A description of the impact the issue is having on your site, using the following definitions:
 - Site Down

All network traffic has ceased, causing a critical impact to your business.

- **Site at Risk**
Primary unit has failed resulting in no redundancy. Site is at risk of going down.
 - **Performance Degraded**
Network traffic is partially functional causing some applications to be unreachable.
 - **General Assistance**
Questions regarding configurations. Troubleshooting non-critical issue or request for product functionality that is not part of the current product feature set.
3. The hours that you are available to work on the issue and any alternative contacts that can work on the issue if you are not available.
 4. Remote access information, if possible.
 - Remote access to your network environment is important, because it is the most effective method for collecting information and troubleshooting technical issues. If you cannot provide remote access, F5 Technical Support will work directly with you to resolve the issue over the phone; however, this method can often be more time consuming and may require file transfers, replication, and additional testing.

tech.out file (qkview)

A tech.out file contains the configuration files that F5 Technical Support most frequently needs when troubleshooting an issue. A tech.out file is produced by the qkview utility and the terms tech.out and qkview may be used interchangeably.

Log files

The tech.out file contains the log files for the last day. If the issue has existed for more than a day, provide all the log files on the system by performing the following procedure:

Packet traces

If the issue involves the network, perform a packet trace while the issue is occurring and provide the packet trace when you open the case.

SSLDUMP

If the issue involves SSL-encrypted packet streams managed by the BIG-IP system, you can use the ssldump utility to examine, decrypt, and decode the SSL-encrypted packets.

UCS archive

If you cannot give F5 Technical Support remote access to your system, you must provide a UCS archive of the current configuration.

Core files

Core files contain the contents of the system memory at the time a crash occurred. If the system has been configured to save core files, they will be located in the `/var/core` directory (BIG-IP 9.3 and later). Provide any existing core files when you open the case.

Conclusion

This document is intended as a study guide for the F5 302 – GTM Specialist: Architect, Set-Up & Deploy exam. This study guide is **not** an all-inclusive document that will guarantee a passing grade on the exam. It is intended to be a living doc and any feedback or material that you feel should be included, to help exam takers better prepare, can be sent to channeleng@f5.com.

Thank you for using this study guide to prepare the 302 – GTM Specialist exam and good luck with your certification goals.

Thanks,

Eric Mitchell
Channel FSE, East US and Federal

