# F5 Networks

## F5CAB5 Exam

### BIG-IP Administration Support and Troubleshooting

**Questions & Answers**
**Demo**

# Version: 4.0

## Question: 1

A BIG-IP Administrator needs to determine why only one pool member is showing connections from the virtual server, resulting in uneven load balancing.

What two reasons would cause uneven load balancing? (Choose two answers)

A. The virtual server is marked down.

B. Monitors have marked down multiple pool members.

C. The pool has a persistence profile configured.

D. All pool members are marked down.

**Answer: B, C**

Explanation:

Uneven load balancing on a BIG-IP system typically occurs when traffic is not distributed evenly across all available pool members. One common reason is that monitors have marked down multiple pool members (Option B). When health monitors fail for specific pool members, BIG-IP automatically removes those members from load-balancing decisions. As a result, traffic is sent only to the remaining healthy member, creating the appearance that load balancing is not functioning correctly. This behavior is expected and aligns with BIG-IP's design to ensure traffic is sent only to healthy resources.

Another frequent cause is the presence of a persistence profile on the pool or virtual server (Option C). Persistence (such as source address or cookie persistence) forces subsequent client connections to be sent to the same pool member for session continuity. While persistence is critical for certain applications, it can override the load-balancing algorithm and cause most or all traffic to be directed to a single pool member, especially during low traffic volumes or testing scenarios.

The other options are incorrect because a virtual server marked down (Option A) would not pass traffic at all, and all pool members marked down (Option D) would result in no connections rather than uneven distribution. This analysis follows standard BIG-IP troubleshooting methodology using pool status, monitor results, and persistence configuration review.

## Question: 2

A BIG-IP Administrator plans to upgrade a BIG-IP device to the latest TMOS version.

Which two tools could the administrator leverage to verify known issues for the target versions? (Choose two answers)

A. F5 End User Diagnostics (EUD)

B. F5 Bug Tracker

C. F5 University

D. F5 iHealth

E. F5 Downloads

**Answer: B, D**

Explanation:

Before upgrading a BIG-IP system to a newer TMOS version, it is critical to review known issues to avoid introducing instability or regressions. F5 Bug Tracker (Option B) is a primary resource for this purpose. It allows administrators to search for documented software defects by TMOS version, module, symptom, or bug ID. Using Bug Tracker, an administrator can identify unresolved issues, fixed bugs, and behavioral changes that may affect their specific deployment, such as traffic handling, high availability, or module-specific functionality. This directly supports proactive troubleshooting and informed upgrade planning.

F5 iHealth (Option D) is another essential tool used during upgrade preparation. iHealth analyzes uploaded UCS or QKView files and correlates the device configuration and software version with F5's known issues database. It provides actionable reports highlighting critical defects, upgrade risks, interoperability concerns, and recommended target versions. iHealth is especially valuable because it contextualizes known issues based on the actual configuration running on the device.

The other options are not appropriate for verifying known software issues. F5 End User Diagnostics (Option A) is a client-side troubleshooting tool, F5 University (Option C) is a training platform, and F5 Downloads (Option E) is primarily used to obtain software images and release notes, not to analyze known defects in depth.

## Question: 3

A BIG-IP Administrator needs to collect HTTP status code and HTTP method for traffic flowing through a virtual server.

Which default profile provides this information? (Choose one answer)

A. Request Adapt

B. HTTP

C. Analytics

D. Statistics

**Answer: C**

Explanation:

To collect application-layer details such as HTTP status codes (200, 404, 500, etc.) and HTTP methods (GET, POST, PUT, DELETE), the BIG-IP system must use a profile designed for traffic visibility and reporting rather than basic traffic handling. The Analytics profile (Option C) is the correct choice because it is specifically designed to collect, store, and present detailed statistics about HTTP and TCP traffic passing through a virtual server.

When an Analytics profile is attached to a virtual server, BIG-IP can record metrics such as HTTP response codes, request methods, URI paths, latency, throughput, and client-side/server-side performance data. These statistics are then accessible through the BIG-IP GUI under Statistics → Analytics, allowing administrators to validate application behavior and troubleshoot performance or functional issues.

The HTTP profile (Option B) enables HTTP protocol awareness and features like header insertion and compression, but it does not provide historical or statistical reporting of HTTP methods and response codes. Request Adapt (Option A) is used for ICAP-based content adaptation, not visibility. Statistics (Option D) is not a standalone profile and does not provide HTTP-level insight.

Therefore, the Analytics profile is the only default profile that fulfills this requirement.

## Question: 4

Which two methods should the BIG-IP Administrator use to troubleshoot a pool member that has been marked DOWN by its health monitor? (Choose two answers)

A. Review the BIG-IP routing table using netstat -rn to show all routes.

B. Enable monitor logging for the pool member that is DOWN.

C. Review the pool and pool-member statistics table for error data.

D. Collect a TCPdump packet capture for the DOWN pool member.

**Answer: B, D**

Explanation:

When a pool member is marked DOWN, it indicates that the configured health monitor is failing. The most effective troubleshooting approach is to focus on the monitor behavior and the actual traffic between BIG-IP and the pool member.

Enabling monitor logging (Option B) is a recommended first step. Monitor logging provides detailed information about why the health check is failing, such as timeouts, connection refusals, incorrect responses, or unexpected status codes. This directly correlates with BIG-IP troubleshooting best practices and allows administrators to confirm whether the failure is due to application behavior, incorrect monitor configuration, or network reachability.

Collecting a TCPdump packet capture (Option D) is also a highly effective method. A packet capture allows the administrator to verify whether the monitor probes are being sent, whether responses are received, and whether packets are being dropped, reset, or malformed. This is especially valuable when diagnosing firewall issues, SSL problems, or application-level failures.

Reviewing pool statistics (Option C) is useful for general monitoring but does not explain why a health monitor is failing. Reviewing the routing table (Option A) is typically unnecessary unless there is evidence of a broader routing issue affecting multiple destinations.

## Question: 5

A BIG-IP Administrator observes the following pool member status message:

Pool /Common/testpool member /Common/10.120.0.5:8090 monitor status down

[/Common/http: up, /Common/http2: down; last error:]

Why is this pool member being marked down? (Choose one answer)

A. The pool member is currently only serving HTTP traffic.

B. The pool member is currently only serving TCP traffic.

C. The pool member is currently only serving UDP traffic.

D. The pool member is currently only serving HTTPS traffic.

**Answer: A**

Explanation:

The pool member is marked DOWN because it is monitored by multiple health monitors, specifically an HTTP monitor and an HTTP/2 monitor. The status message clearly shows that the HTTP monitor is UP, while the HTTP/2 monitor is DOWN. In BIG-IP, when multiple monitors are assigned to a pool member, the default behavior is AND logic, meaning all assigned monitors must succeed for the pool member to be considered healthy.

In this scenario, the server is responding successfully to standard HTTP (likely HTTP/1.1) requests but does not support or respond correctly to HTTP/2 requests. As a result, the HTTP/2 monitor fails, which causes the overall monitor status to be DOWN, even though HTTP traffic itself is working.

This behavior is expected and documented in BIG-IP monitoring logic. Unless the monitor rule is explicitly changed to “at least one of”, a single failing monitor will mark the pool member down. Therefore, the correct conclusion is that the pool member is only serving HTTP traffic, not HTTP/2.

The resolution would be to either remove the HTTP/2 monitor, correct the application to support HTTP/2, or adjust the monitor rule to match the intended health-check logic.