PassQueen
a.k.a dumpsqueen.com

BIG-IP Administration Support and Troubleshooting () exam

F5 F5CAB5

Version Demo

Total Demo Questions: 8

Total Premium Questions: 85

Buy Premium PDF

https://passqueen.com

support@passqueen.com

QUESTION NO: 1

A BIG-IP Administrator plans a TMOS upgrade and wants to retain both a recoverable configuration copy and a diagnostic baseline of the device before making changes.

Which two actions should the administrator perform? (Choose two answers)

A. Create a UCS archive.

B. Generate a QKView file.

C. Clear the LTM and system logs.

D. Install the target image directly to the active boot volume.

E. Reset all virtual-server statistics.

ANSWER: A B

Explanation:

Creating a UCS archive preserves the BIG-IP configuration and associated system files needed to restore the configuration if required. Generating a QKView captures diagnostic information, including logs, system status, and configuration details, that can be used to compare the pre-upgrade state or assist F5 Support if an issue occurs.

Saving the running configuration alone does not create a separate recovery archive. Clearing logs removes useful diagnostic history, and installing an image directly to the active boot volume does not provide a safe upgrade workflow.

QUESTION NO: 2

A BIG-IP Administrator manually updates the configuration files during a maintenance window and wants to validate the configuration syntax before loading it into the running configuration.

Which command should the administrator use?

A. `tmsh load sys config verify`

B. `tmsh save sys config`

C. `tmsh load sys config`

D. `bigstart restart mcpd`

ANSWER: A

Explanation:

`tmsh load sys config verify` validates the BIG-IP configuration files without loading the configuration into the running system. This allows the administrator to identify syntax or reference errors before applying the edited configuration.

`tmsh save sys config` writes the current running configuration to disk, while `tmsh load sys config` loads the configuration and can affect the running system. Restarting MCPD is not a configuration-validation method.

QUESTION NO: 3

Users report intermittent TCP resets through a production virtual server. The issue is occurring now, and the BIG-IP Administrator is preparing information for F5 Support.

Which two data collections provide the most useful initial evidence? (Choose two answers)

A. Generate a QKView file after the issue is observed.

B. Collect a targeted tcpdump packet capture while the resets occur.

C. Create only a UCS archive.

D. Restart TMM before collecting diagnostic information.

E. Clear all virtual-server statistics before collecting a packet capture.

ANSWER: A B

Explanation:

A current QKView provides F5 Support with a broad diagnostic snapshot, including system configuration, log files, platform information, and operational status. A targeted packet capture collected while the issue occurs can show the TCP reset packet, the traffic direction, and which endpoint transmitted it.

A UCS archive is primarily a configuration backup and does not replace diagnostic logs and runtime information in a QKView. Clearing counters or restarting TMM before gathering evidence can remove useful information and can disrupt production traffic.

QUESTION NO: 4

Which two methods should the BIG-IP Administrator troubleshoot a Pool-member that's been marked "DOWN" by its Health Monitor? (Pick the 2 correct responses below)

A. Enable Monitor Logging for the Pool-member that's "DOWN".

B. Review the BIG-IP's routing table using "netstat -rn" to show all routes.

C. Collect a TCPdump packet capture for the "DOWN" Pool-member.

D. Review the Pool & Pool-member Statistics table for error data.

ANSWER: A C

Explanation:

Enable Monitor Logging for the Pool-member that's "DOWN" and collect a TCPdump packet capture for the "DOWN" Pool-member because these methods provide the monitor-level evidence needed to identify why BIG-IP declared the member unavailable. Monitor logging records the health monitor transaction details, including the request sent by BIG-IP and the response received, or the absence of an expected response. This lets the administrator validate monitor settings such as the destination, request string, response string, interval, timeout, and application-specific behavior against what the server actually returns.

A packet capture complements monitor logging by showing the network exchange on the relevant VLAN. It can confirm whether monitor packets leave the BIG-IP system, reach the pool member, and receive an expected response. It also exposes conditions that can prevent a successful monitor result, such as TCP resets, failed handshakes, ICMP errors, asymmetric routing, or a missing return packet. Together, these techniques distinguish application-response failures from transport and network-path failures, providing the evidence required to correct the monitor configuration or the underlying service. F5 documents monitor configuration and behavior in the F5 health monitor documentation and provides packet-capture guidance in F5 TCPDUMP guidance.

QUESTION NO: 5

A BIG-IP Administrator configured the following virtual server to pass traffic on all addresses and ports. After configuration is completed, the BIG-IP Administrator notices that the virtual server is unable to pass traffic.

ltm virtual forwarding_any_vs {

destination 0.0.0.0:any

ip-forward

mask 255.255.255.255

profiles {

fastL4 { }

}

serverssl-use-sni disabled

source 0.0.0.0/0

translate-address disabled

translate-port disabled

}

Which part of the configuration is the cause of the issue? (Choose one answer)

A. Incorrect destination configured

B. Incorrect mask 255.255.255.255

C. Incorrect translate-address configured

ANSWER: B

Explanation:

Incorrect mask 255.255.255.255 is the cause of the problem. An IP-forwarding virtual server intended to accept traffic for every IPv4 destination must use a wildcard destination range. Although the destination is displayed as `0.0.0.0:any`, the mask determines how much of that destination address is matched. A `255.255.255.255` mask is a host mask, so the virtual server matches only the single address `0.0.0.0`, rather than all destination addresses. Because ordinary forwarded packets are addressed to real routable destination addresses, they do not match this virtual server and therefore cannot be handled by its IP-forwarding configuration.

For an all-addresses forwarding virtual server, the destination must be paired with a wildcard mask of `0.0.0.0`. This creates the effective match of `0.0.0.0/0`, allowing the BIG-IP system to select the virtual server regardless of the packet's destination IPv4 address. The `ip-forward` setting then directs BIG-IP to route matching traffic rather than load-balance it to a pool. F5 documents the virtual server `destination` and `mask` properties in the tmsh ltm virtual reference.

QUESTION NO: 6

A web application uses HTTP keep-alive connections. The virtual server has an HTTP profile and a Round Robin pool, but no persistence profile. During testing, one client sends many HTTP requests over a single TCP connection and all requests are processed by the same pool member.

What should the BIG-IP Administrator configure if individual HTTP requests must be eligible for distribution to different pool members?

A. Configure a OneConnect profile.

B. Configure source-address persistence.

C. Change the pool method to Least Connections (Node).

D. Configure a Server SSL profile.

ANSWER: A

Explanation:

Configure a **OneConnect profile** on the virtual server. Without OneConnect, BIG-IP generally maintains the server-side connection associated with the client-side TCP connection. Multiple HTTP requests sent over one keep-alive client connection therefore continue to use the initially selected pool member.

A OneConnect profile allows BIG-IP to reuse server-side connections and make load-balancing decisions at HTTP request boundaries, subject to the configured OneConnect behavior. Changing the pool to Least Connections does not alter the fact that one persistent TCP connection has already selected a member. Source-address persistence would further increase affinity rather than distribute requests.

QUESTION NO: 7

Refer to the exhibit.

The image shows the status of a virtual server namedapplication_vsin the BIG-IP Configuration Utility.

What is the cause of the status shown? (Choose two answers)

A. Pool member(s) administratively disabled

B. Pool member(s) forced offline

C. Node(s) administratively disabled

D. Virtual Server administratively disabled

ANSWER: A B C

Explanation:

An enabled virtual server can display an offline status when BIG-IP has no usable pool member to which it can direct new connections. “Pool member(s) administratively disabled” causes that condition when every eligible member is disabled: BIG-IP stops selecting disabled members for new traffic, although existing connections may be allowed to drain. “Pool member(s) forced offline” also produces the condition when it affects all eligible members. Forced offline removes a member from service immediately and prevents it from receiving either new connections or continued traffic. Finally, “Node(s) administratively disabled” can produce the same virtual-server status because a pool member depends on the availability and administrative state of its parent node. Disabling a node makes its associated members unavailable for load balancing. Therefore, if these actions leave no available members in the virtual server’s default pool, the virtual server remains administratively enabled but is operationally offline. F5 documents pool-member administrative states, including disabled and forced-offline behavior, in the tmsh ltm pool members reference. The relationship between nodes, members, pools, and virtual-server availability is also covered in the BIG-IP Local Traffic Manager basics documentation.

QUESTION NO: 8

Pool /Common/testpool member /Common/10.120.0.5:8090 monitor status down. [/Common/http: up, /Common/http2: down; last error:] [was up for lhr:0min:43sec]

Why is this pool member being marked down?

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

B. The pool member is currently only serving HTTP2 traffic.

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

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

ANSWER: C

Explanation:

The pool member is currently only serving HTTP traffic. The status output shows that the assigned `/Common/http` monitor is up, confirming that the member is successfully responding to the HTTP health check. At the same time, the `/Common/http2` monitor is down, so the member is not passing the HTTP/2-specific health check. When multiple monitors are applied through a monitor rule with the normal all-monitors-required behavior, BIG-IP considers the member available only when every required monitor succeeds. Consequently, the failed HTTP/2 monitor causes the aggregate monitor status, and therefore the pool-member status, to be marked down despite the successful HTTP monitor result.

This output does not merely indicate a generic connectivity issue: it distinguishes the successful HTTP monitor from the failed HTTP/2 monitor. The evidence supports that the service is answering the standard HTTP probe but is not meeting the configured HTTP/2 monitor requirements. BIG-IP monitor rules can also be configured with an availability requirement that changes how many monitors must succeed; absent such a rule, all assigned monitors must report success for the member to be available. See the F5 ltm pool tmsh reference for monitor-rule behavior and the HTTP/2 monitor tmsh reference for HTTP/2 monitor configuration.