Vendor: F5

Exam Code: F5CAB2

Exam Name: BIG-IP Administration Data Plane Concepts

Version: DEMO

QUESTION 1
A BIG-IP system receives a client connection destined to 1.0.0.10:8080. Multiple virtual servers are configured on the system. Which virtual server will process the connection? (Choose one answer)

A. A forwarding virtual server configured with 0.0.0.0:any
B. A forwarding virtual server configured with 1.0.0.10:any (port 0)
C. A virtual server configured with destination 1.0.0.10:8080 and is available (green)
D. A virtual server configured with 0.0.0.0:8080

Answer: C
Explanation:
BIG-IP uses a virtual server matching and precedence algorithm to determine which virtual server processes an incoming connection. This decision is made entirely in the data plane and is based on how specifically a virtual server matches the destination IP address and port.

BIG-IP Virtual Server Selection Rules (Simplified):

When multiple virtual servers could match a packet, BIG-IP selects the most specific match, using the following precedence:

Exact IP address and exact port

Exact IP address with wildcard port (port 0 / any)

Wildcard IP address with exact port

Wildcard IP address and wildcard port

Applying the Rules to This Scenario:

Incoming traffic destination: 1.0.0.10:8080

Option C: 1.0.0.10:8080

Exact IP match

Exact port match

Highest possible specificity

If the virtual server is available (green), it wins the match

QUESTION 2
Refer to the exhibit above. A BIG-IP pool is configured with Priority Group Activation = Less than 2 available members. The pool members have different priority groups and availability states. Which pool members are receiving traffic? (Choose one answer)

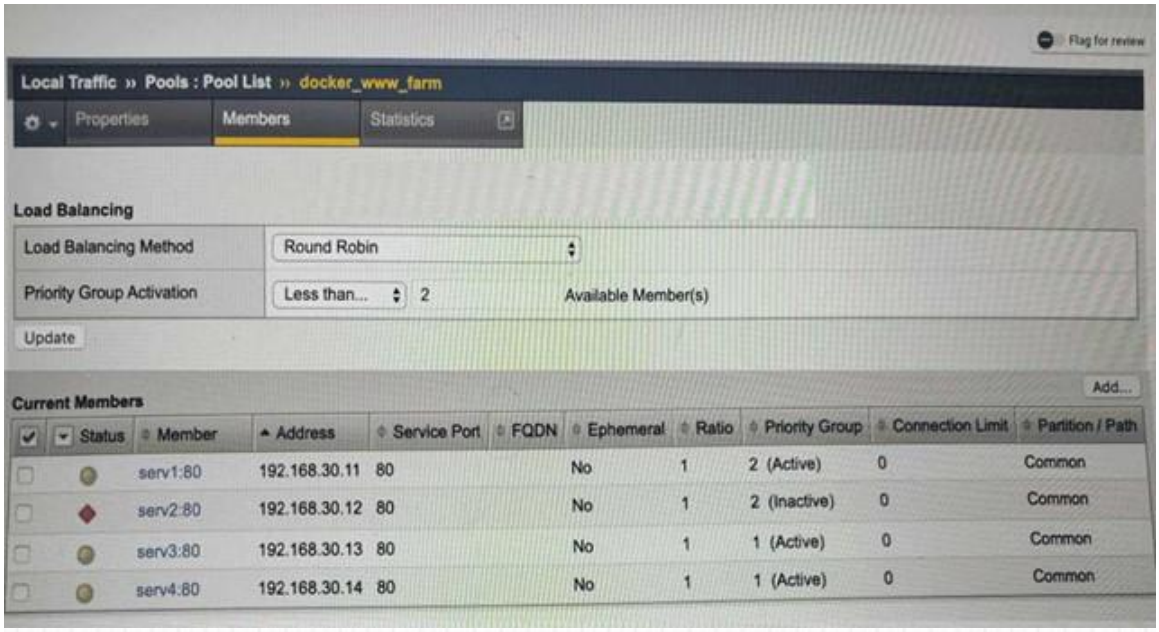
Flag for review
Local Traffic » Pools : Pool List » docker_www_farm
Properties
Members
Statistics
Load Balancing
Load Balancing Method
Round Robin
Priority Group Activation
Less than...
2
Available Member(s)
Update
Current Members
Add...
Status
Member
Address
Service Port
FQDN
Ephemeral
Ratio
Priority Group
Connection Limit
Partition / Path
serv1:80 192.168.30.11 80 No 1 2 (Active) 0 Common
serv2:80 192.168.30.12 80 No 1 2 (Inactive) 0 Common
serv3:80 192.168.30.13 80 No 1 1 (Active) 0 Common
serv4:80 192.168.30.14 80 No 1 1 (Active) 0 Common

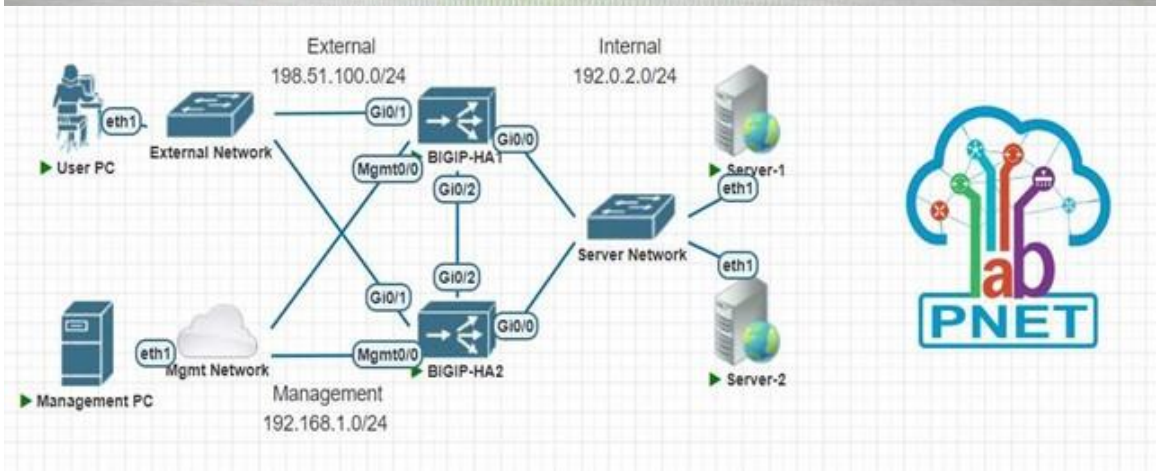
External
198.51.100.0/24
Internal
192.0.2.0/24
eth1
User PC
External Network
Gi0/1
Mgmt0/0
BIGIP-HA1
Gi0/0
Gi0/2
Server-1
eth1
Server Network
eth1
Server-2
Gi0/2
Gi0/1
Gi0/0
eth1
Mgmt Network
Mgmt0/0
BIGIP-HA2
Management PC
Management
192.168.1.0/24
PNET

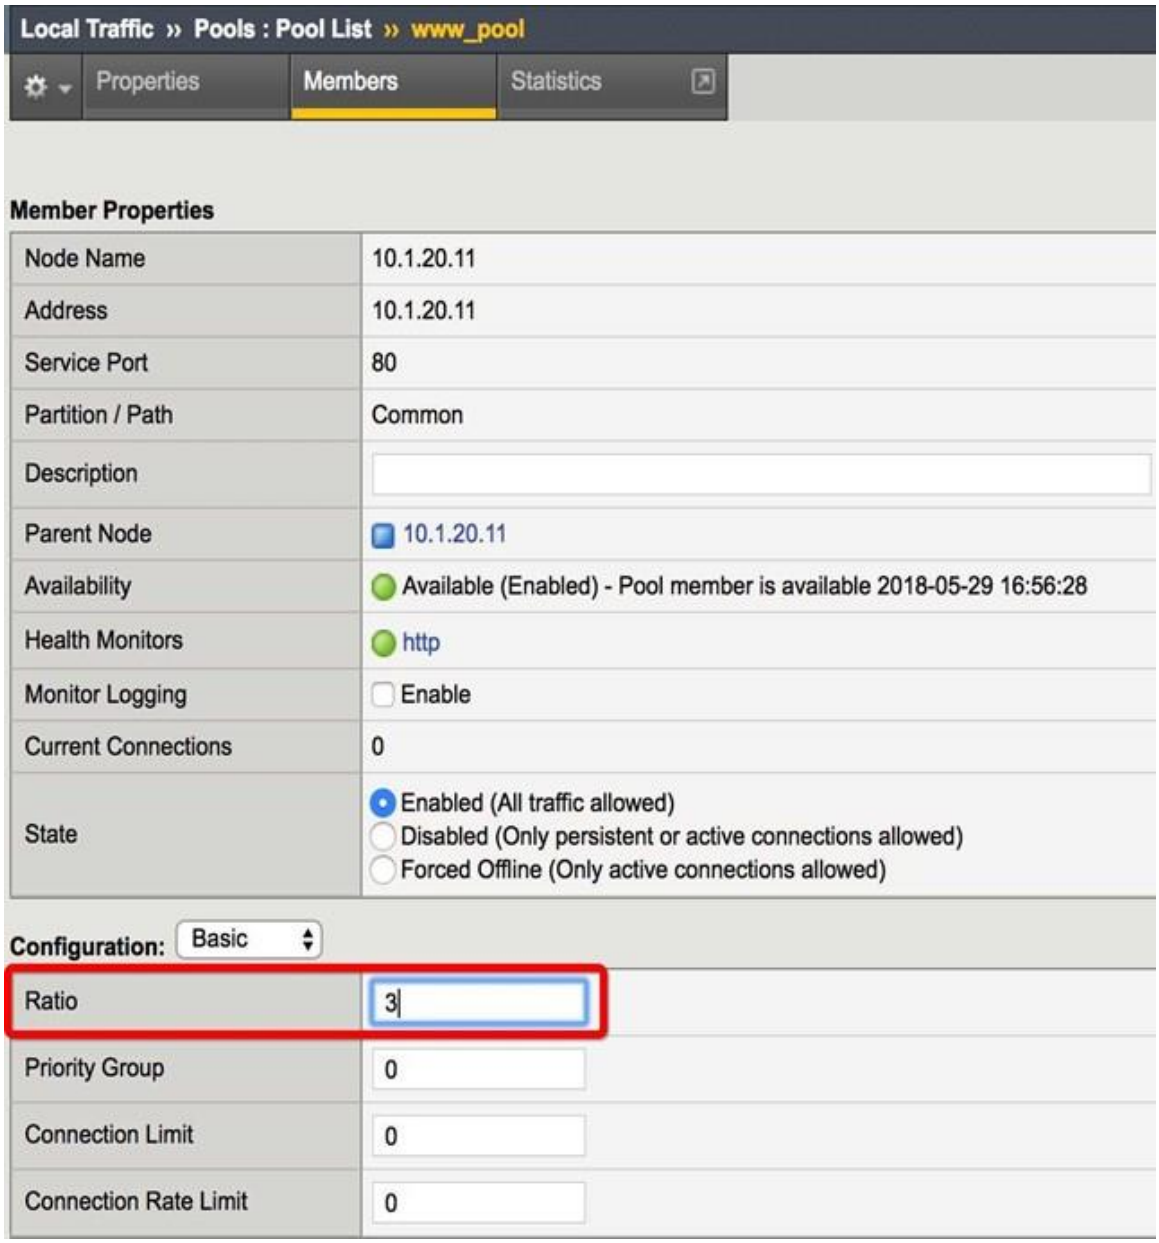
Local Traffic » Pools : Pool List » www_pool
Properties
Members
Statistics
Member Properties
Node Name
10.1.20.11
Address
10.1.20.11
Service Port
80
Partition / Path
Common
Description
Parent Node
10.1.20.11
Availability
Available (Enabled) - Pool member is available 2018-05-29 16:56:28
Health Monitors
http
Monitor Logging
Enable
Current Connections
0
State
Enabled (All traffic allowed)
Disabled (Only persistent or active connections allowed)
Forced Offline (Only active connections allowed)
Configuration:
Basic
Ratio
3
Priority Group
0
Connection Limit
0
Connection Rate Limit
0

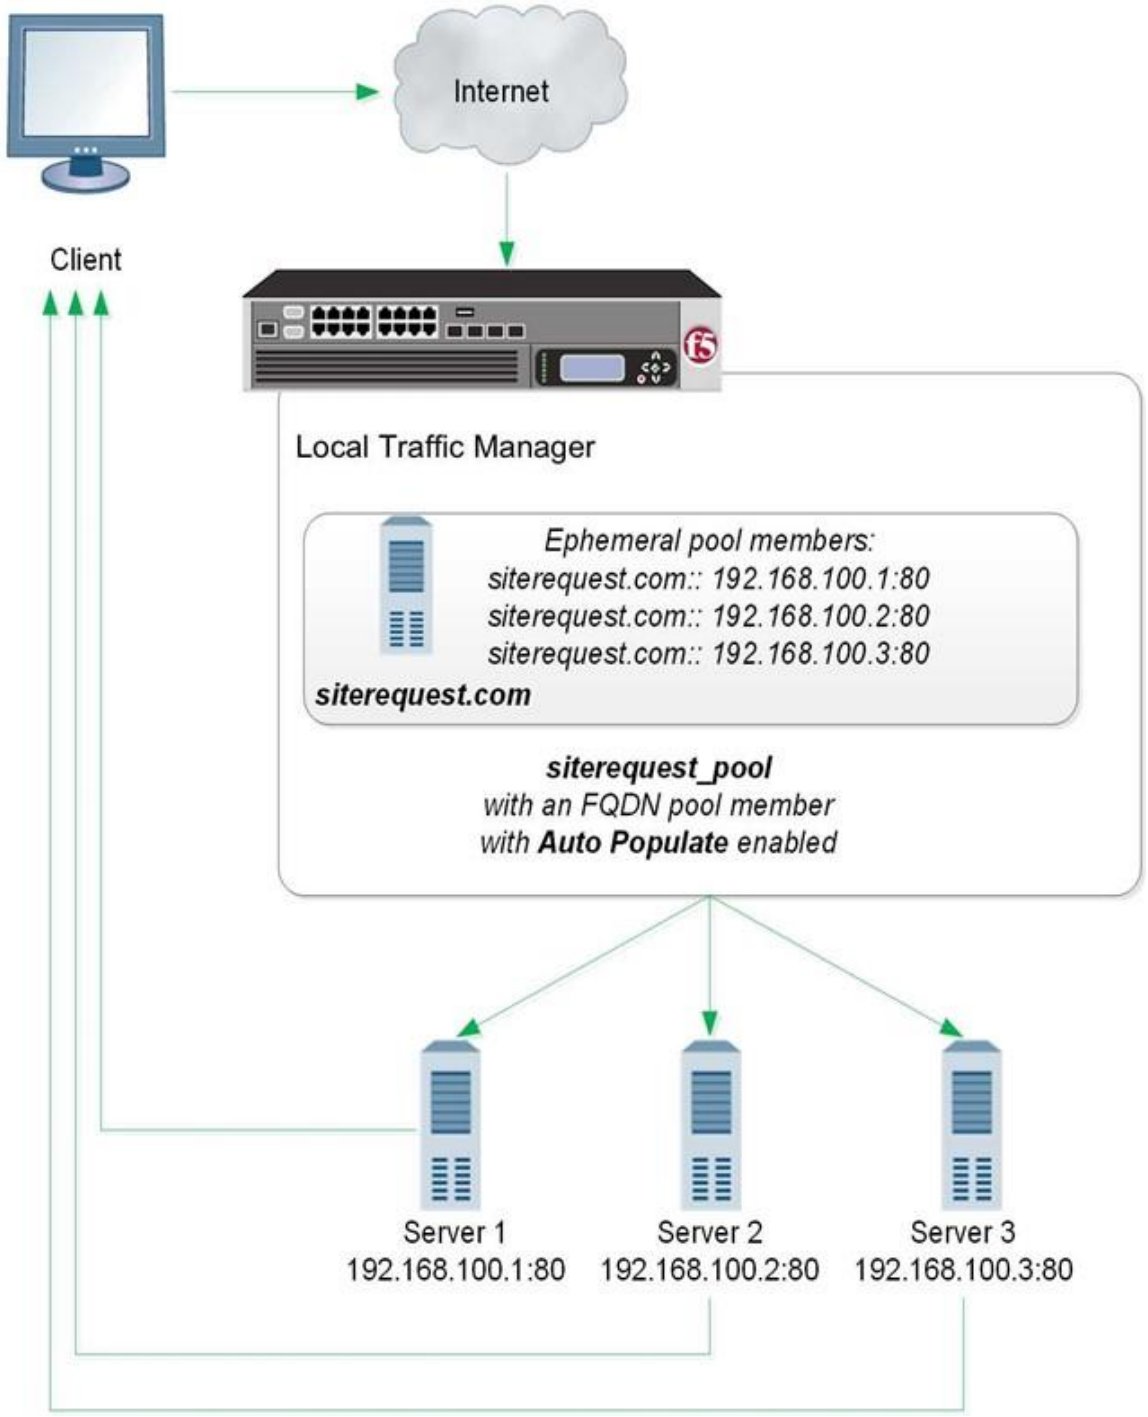

A. serv1
B. serv1, serv3
C. serv1, serv2, serv3, serv4
D. serv1, serv3, serv4

Answer: D
Explanation:
This question tests understanding of Priority Group Activation (PGA) and how BIG-IP determines

which pool members are eligible to receive traffic.

Key BIG-IP Priority Group Concepts:

Higher priority group numbers = higher priority

BIG-IP will only send traffic to the highest priority group that meets the Priority Group Activation condition

Lower priority groups are activated only when the condition is met

Only available (green) members count toward the activation threshold

Configuration from the Exhibit:

Priority Group Activation: Less than 2 available members

Pool Members and Status:

Pool Member Priority Group Status

serv1 2 Active (available)

serv2 2 Inactive (down)

serv3 1 Active (available)

serv4 1 Active (available)

Step-by-Step Traffic Decision:

BIG-IP first evaluates the highest priority group (Priority Group 2)

Priority Group 2 has:

serv1 -> available

serv2 -> unavailable

Total available members = 1

Activation rule is Less than 2 available members

Condition is true (1 < 2)

BIG-IP activates the next lower priority group (Priority Group 1)

Traffic is now sent to:

serv1 (Priority Group 2)

serv3 and serv4 (Priority Group 1)

Final Result:

Traffic is distributed to serv1, serv3, and serv4

QUESTION 3
The network architecture for a BIG-IP consists of an external VLAN and an internal VLAN with two interfaces connected to the upstream switch. The design requires fault tolerance in the case that one of the interfaces is down. Which deployment architecture meets these requirements? (Choose one answer)

A. One network trunk with both VLANs and LACP enabled, and both VLANs configured as untagged
B. Two network trunks each with one VLAN and LACP enabled, and both VLANs configured as tagged
C. Two network trunks each with one VLAN and LACP disabled, and one VLAN configured as tagged and one VLAN configured as untagged
D. One network trunk with both VLANs and LACP enabled, and both VLANs configured as tagged

Answer: D
Explanation:
To meet the requirement of fault tolerance when one interface goes down, BIG-IP must use link aggregation so that loss of a single physical link does not isolate the VLAN(s).

How the objects relate (data plane view)

Interfaces = physical links.

Trunk (LACP) = bundles multiple interfaces into one logical link that provides redundancy (and possibly bandwidth aggregation).

VLANs are assigned to interfaces or trunks. If you need multiple VLANs on the same trunk, they must use 802.1Q tagging (because you can only have one untagged VLAN per interface/trunk).

Self IPs are then placed on the VLANs to provide BIG-IP presence and routing/ARP functions, but self IPs are not what provides link resiliency--the trunk does.

Why Option D is correct

You have two physical interfaces and you want resiliency if one fails put both interfaces into one trunk with LACP enabled.

You need both external and internal VLANs on those same two links both VLANs should be configured as tagged on that trunk, so they can coexist on the same aggregated link.

If either physical interface fails, the trunk remains up via the remaining interface, keeping both VLANs operational.

QUESTION 4
What should a BIG-IP Administrator configure to minimize impact during a failover? (Choose one answer)

A. External monitors
B. Clone pool
C. OneConnect profile
D. MAC masquerading

Answer: D

Explanation:
In BIG-IP high availability (HA) deployments, one of the primary causes of traffic disruption during failover is Layer 2 and Layer 3 relearning by upstream network devices (switches and routers). When traffic groups move from the Active device to the Standby device, the network must quickly associate the IP addresses with the new device.

Why MAC Masquerading Minimizes Failover Impact:

MAC masquerading allows a traffic group to use a floating, shared MAC address for its Self IPs. This MAC address moves with the traffic group during failover.

Key benefits:

The MAC address does not change when failover occurs

Upstream switches do not need to relearn ARP entries

Traffic resumes almost immediately after failover

Dramatically reduces packet loss and connection interruption

From BIG-IP Administration Data Plane Concepts:

MAC masquerade is specifically designed to provide fast failover

It is a best practice for HA pairs, especially in environments sensitive to latency and connection loss

QUESTION 5
Refer to the exhibit. The BIG-IP Administrator needs to avoid overloading any of the pool members with connections when they become active. What should the BIG-IP Administrator configure to meet this requirement? (Choose one answer)

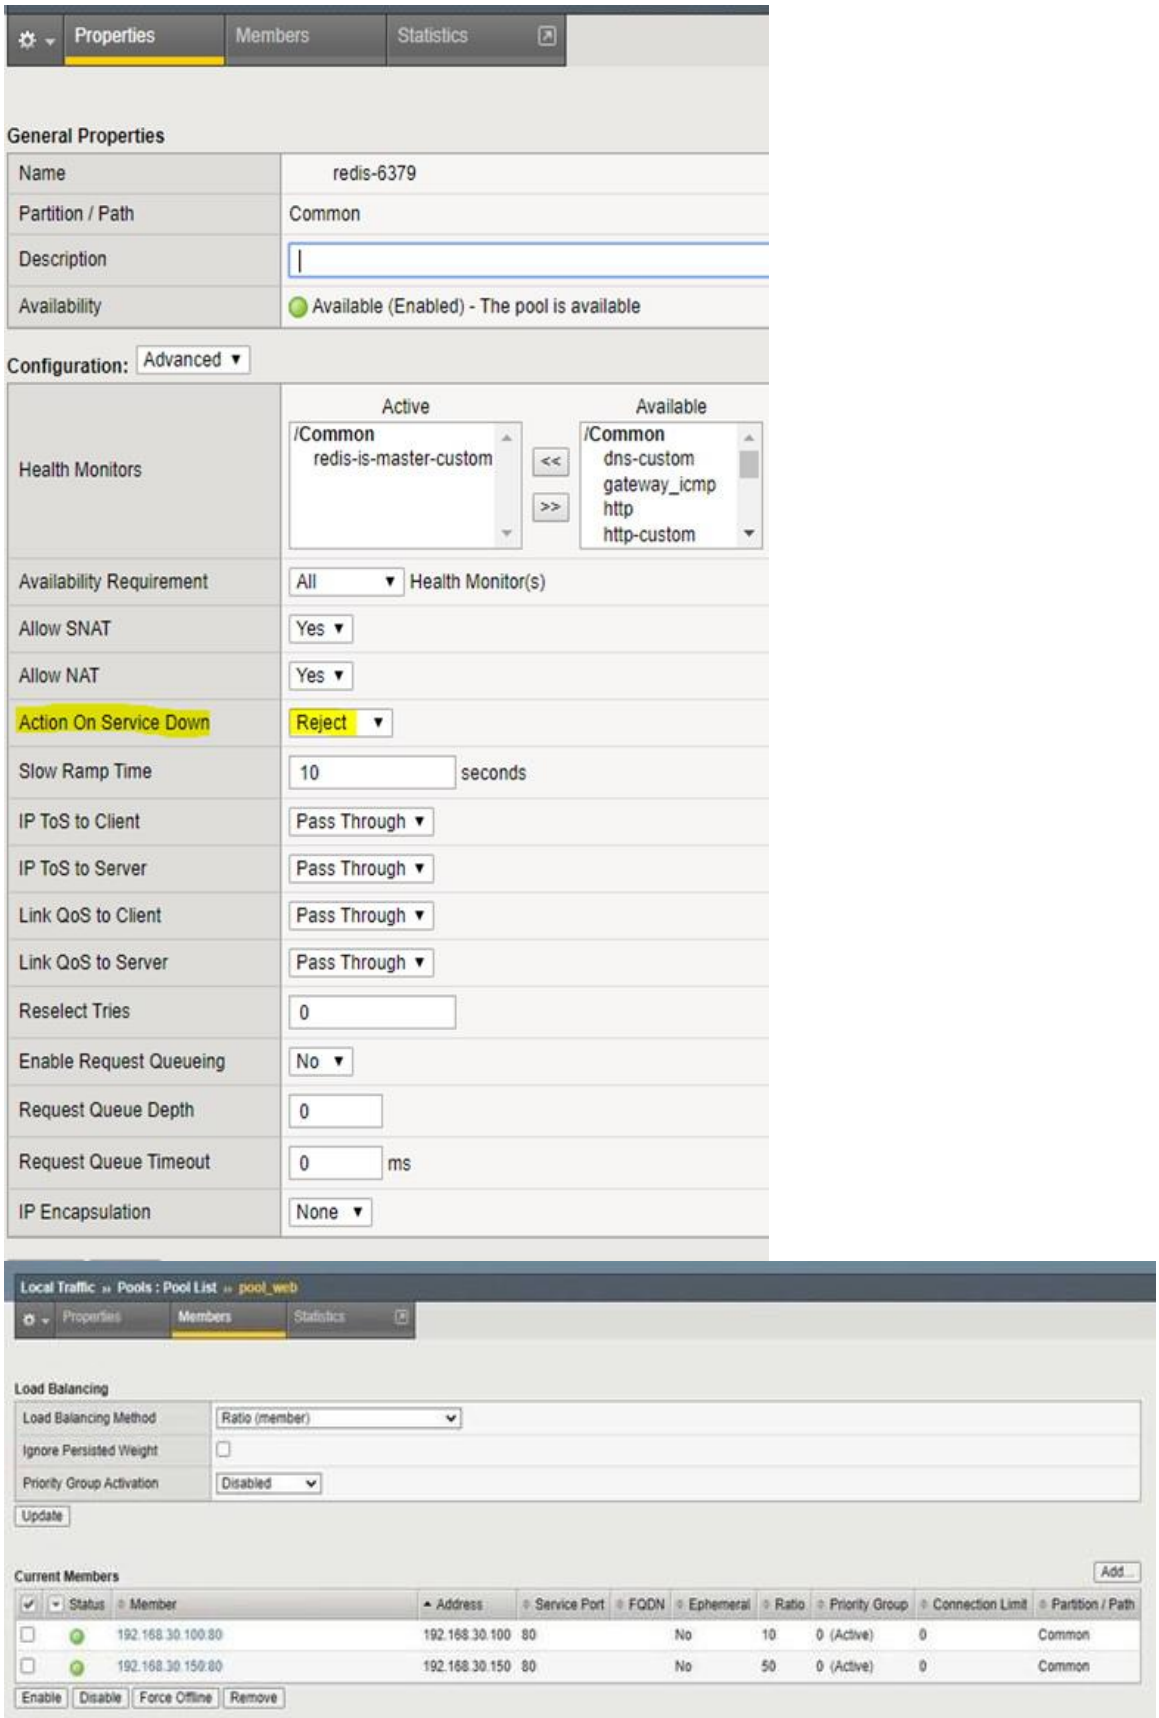
Properties
Members
Statistics
General Properties
Name
redis-6379
Partition / Path
Common
Description
Availability
Available (Enabled) - The pool is available
Configuration: Advanced
Active
Available
/Common
redis-is-master-custom
/Common
dns-custom
gateway_icmp
http
http-custom
Health Monitors
Availability Requirement
All
Health Monitor(s)
Allow SNAT
Yes
Allow NAT
Yes
Action On Service Down
Reject
Slow Ramp Time
10
seconds
IP ToS to Client
Pass Through
IP ToS to Server
Pass Through
Link QoS to Client
Pass Through
Link QoS to Server
Pass Through
Reselect Tries
0
Enable Request Queueing
No
Request Queue Depth
0
Request Queue Timeout
0
ms
IP Encapsulation
None
Local Traffic » Pools : Pool List » pool_web
Properties
Members
Statistics
Load Balancing
Load Balancing Method
Ratio (member)
Ignore Persisted Weight
Priority Group Activation
Disabled
Update
Current Members
Add...
Status
Member
Address
Service Port
FQDN
Ephemeral
Ratio
Priority Group
Connection Limit
Partition / Path
192.168.30.100:80
192.168.30.100
80
No
10
0 (Active)
0
Common
192.168.30.150:80
192.168.30.150
80
No
50
0 (Active)
0
Common
Enable
Disable
Force Offline
Remove

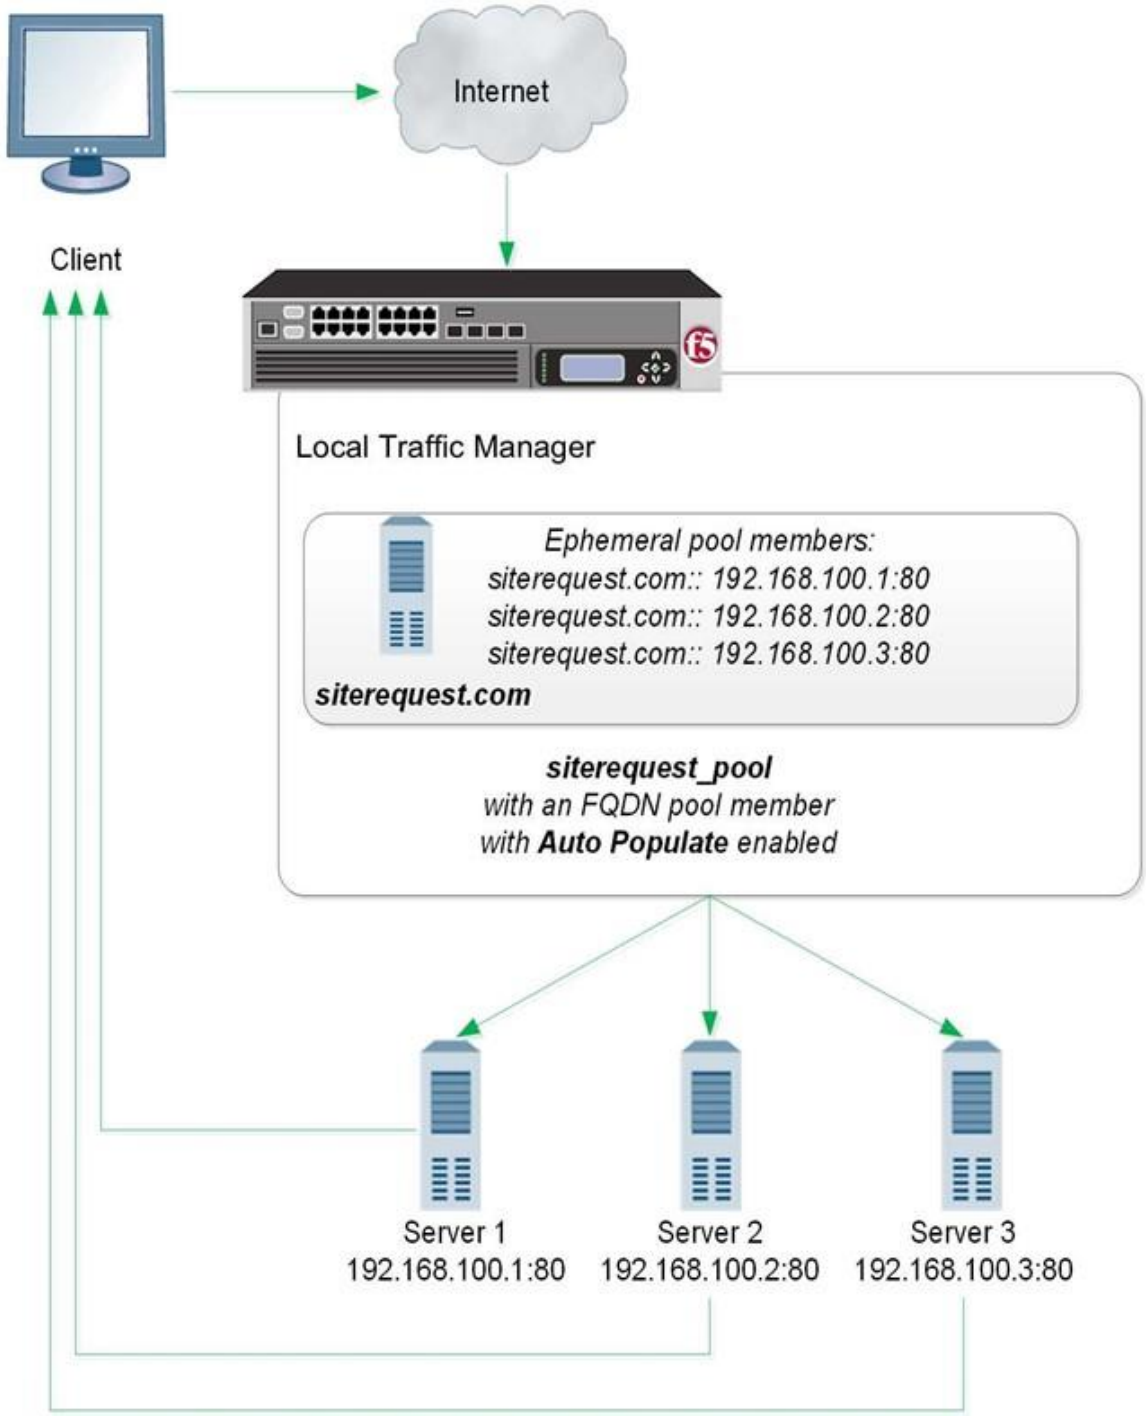

A. Slow Ramp Time to the Pool
B. Different Ratio for each member
C. Action On Service Down to Reselect
D. Same Priority Group to each member

Answer: A
Explanation:
This question focuses on connection behavior when pool members transition from down to up,

which is a classic data plane consideration in BIG-IP environments.

What problem is being solved?

When a pool member:

Recovers from a failure

Is enabled after maintenance

Transitions from inactive to active

...it can suddenly receive a large burst of new connections, especially when using load-balancing methods such as Least Connections. This sudden surge can overload the server.

Why Slow Ramp Time is the correct solution:

Slow Ramp Time is a pool-level setting that:

Gradually increases the number of connections sent to a newly available pool member

Prevents sudden spikes in traffic

Allows the server to warm up (application cache, JVM, DB connections, etc.)

From BIG-IP Administration Data Plane Concepts:

Slow Ramp Time controls the rate at which BIG-IP increases load to a pool member that has just become available

During the ramp period, BIG-IP artificially increases the member's connection count, making it appear "busier" and therefore less attractive for new connections

This directly satisfies the requirement to avoid overloading pool members when they become active.

QUESTION 6
A BIG-IP Administrator configures remote authentication and needs to ensure that users can still log in even when the remote authentication server is unavailable. Which action should the BIG-IP Administrator take in the remote authentication configuration to meet this requirement? (Choose one answer)

A. Configure a second remote user directory
B. Configure a remote role group
C. Enable the Fallback to Local option
D. Set partition access to All

Answer: C
Explanation:
Although remote authentication (LDAP, RADIUS, TACACS+) is a control-plane / management-plane feature, it directly affects availability and resiliency of administrative access, which is a critical operational HA consideration.

How BIG-IP Remote Authentication Works:

BIG-IP can authenticate administrators against:

LDAP

RADIUS

TACACS+

When remote authentication is enabled, BIG-IP by default relies on the remote server for user authentication

If the remote authentication server becomes unreachable, administrators may be locked out unless fallback is configured

Why "Fallback to Local" Is Required:

The Fallback to Local option allows BIG-IP to:

Attempt authentication against the remote authentication server first

If the remote server is unreachable or unavailable, fall back to:

Local BIG-IP user accounts (admin, or other locally defined users)

This ensures:

Continuous administrative access

Safe recovery during:

Network outages

Authentication server failures

Maintenance windows

This behavior is explicitly recommended as a best practice in BIG-IP administration to avoid loss of management access.

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

★ More than **99,900** Satisfied Customers Worldwide.

★ Average **99.9%** Success Rate.

★ **Free Update** to match latest and real exam scenarios.

★ **Instant Download** Access! No Setup required.

★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.

★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.

★ **100%** Guaranteed Success or **100%** Money Back Guarantee.

★ **Fast**, helpful support **24x7**.

View list of all certification exams: http://www.lead2pass.com/all-products.html

10% Discount Coupon Code: **ASTR14**