

DUMPSBOSS.

F5 Cloud Solutions

F5 402

Version Demo

Total Demo Questions: 8

Total Premium Questions: 87

Buy Premium PDF

<https://dumpsboss.co>

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Understand Cloud Architecture	50
Topic 2, Understand Cloud Networking	6
Topic 3, Understand Cloud Security	5
Topic 4, Understand F5 Cloud Solutions	13
Topic 5, Deploy F5 Cloud Solutions	12
Topic 6, Mix Questions	1
Total	87

QUESTION NO: 1

Scenario: A development team deploys containerized applications on Kubernetes and wants BIG-IP to update application-delivery configuration automatically as Services and Ingress resources change. The team plans to use F5 Container Ingress Services (CIS).

Which TWO requirements should be addressed?

(Select TWO)

Choose:

- A. Grant CIS a Kubernetes service account with the required RBAC permissions
- B. Provide network connectivity from BIG-IP to the selected application endpoints
- C. Expose the Kubernetes API publicly to allow CIS access
- D. Manually recreate BIG-IP objects after every Kubernetes update

ANSWER: A B

Explanation:

Grant CIS a Kubernetes service account with the required RBAC permissions so it can watch the Kubernetes resources that drive BIG-IP configuration. Without appropriate API access, CIS cannot reliably detect Service, Ingress, or related resource changes.

Provide network connectivity from BIG-IP to the selected application endpoints so configured pool members can receive traffic and be monitored. The design must account for the Kubernetes networking model, address reachability, routing, and any required firewall rules.

Exposing the Kubernetes API publicly is not required and can increase risk. Manually recreating BIG-IP objects after every Kubernetes update defeats the automated controller-based integration that CIS is intended to provide.

QUESTION NO: 2

Scenario: During a cloud deployment, you notice that application performance is not meeting expectations. What steps can you take to optimize instance sizing and location?

Choose one:

- A. Increase instance sizes indiscriminately to handle more traffic.
- B. Ignore performance issues and hope they resolve on their own.
- C. Analyze application traffic patterns and latency requirements, then right-size instances and place them in appropriate regions.
- D. Move all instances to a single region to simplify management.

ANSWER: C

Explanation:

Analyze application traffic patterns and latency requirements, then right-size instances and place them in appropriate regions. Effective cloud performance optimization starts with measurement: review request volume, throughput, concurrent connections, CPU and memory consumption, network utilization, response-time metrics, and peak-versus-baseline demand. Those observations allow the deployment team to select instance types and counts that provide the required capacity without overprovisioning. Location is equally important because placing application components and traffic-processing services near users, dependent services, and data sources can reduce round-trip latency and avoid unnecessary cross-region network paths. The selected region or regions should also meet availability, compliance, and service-capability requirements. This is an iterative process: establish performance targets, monitor the deployed environment, adjust sizing or

scaling policies as demand changes, and validate the resulting user experience. F5 cloud deployments benefit from the same disciplined approach, since the performance of traffic-management and security services depends on both available instance resources and the network path to protected applications. AWS describes selecting resources based on measured workload requirements in its [Performance Efficiency Pillar](#), while F5 provides deployment guidance and sizing-related documentation through [F5 CloudDocs](#).

QUESTION NO: 3

Scenario: A cloud-hosted three-tier application must accept HTTPS traffic from the internet while keeping application servers and databases unreachable from public networks. BIG-IP VE must enforce application-delivery policies before requests reach the application tier.

Which deployment approach best meets this requirement?

Choose one:

- A. Place BIG-IP VE between an internet-facing subnet and private application subnets
- B. Place application servers directly in a public subnet with public addresses
- C. Allow internet clients to connect directly to the database tier
- D. Deploy BIG-IP VE only on a management network outside the traffic path

ANSWER: A

Explanation:

The best approach is to **place BIG-IP VE in the controlled path between an internet-facing subnet and private application subnets**. BIG-IP can receive client traffic on its external side, apply SSL/TLS, load-balancing, and security policies, and forward approved traffic to application servers that do not have public addresses.

Placing application servers directly in a public subnet unnecessarily exposes them to the internet. Allowing the database tier to receive client connections bypasses the application architecture and violates segmentation principles. Deploying BIG-IP only on a management network does not place it in the application traffic path.

QUESTION NO: 4

Scenario: An organization wants to make an application portable between two cloud providers. The application currently stores user sessions on individual application instances and keeps configuration files locally on each virtual machine.

Which TWO changes most directly improve application mobility?

(Select TWO)

Choose:

- A. Externalize session state from individual application instances.
- B. Store application configuration outside local virtual-machine files.
- C. Increase the size of each virtual machine.
- D. Use provider-specific infrastructure interfaces throughout the application.

ANSWER: A B

Explanation:

Externalizing session state enables requests to be served by instances in either cloud environment without requiring a user to remain connected to the original server. This is particularly important when traffic is distributed across environments or workloads must be restarted elsewhere.

Externalizing configuration into a managed configuration system or deployment process removes dependence on manually maintained local files. It enables the same application configuration to be reproduced consistently in another cloud environment.

Increasing the size of each virtual machine does not improve portability. Tightly coupling the application to provider-specific infrastructure interfaces increases migration effort and reduces mobility.

QUESTION NO: 5

Scenario: A CI/CD pipeline deploys F5 and cloud application services through APIs. The security team requires the pipeline to make only the required changes and does not want long-lived administrator credentials stored in pipeline configuration.

Which TWO IAM practices should be implemented?

(Select TWO)

Choose:

- A. Use a dedicated workload identity with least-privilege permissions.
- B. Use short-lived credentials or role-based authentication for the pipeline.
- C. Use one shared administrator account for all automation workflows.
- D. Disable audit logging to reduce pipeline activity.

ANSWER: A B

Explanation:

A dedicated workload identity with narrowly scoped permissions applies least privilege to the automation process. The identity should be permitted to perform only the required deployment actions against the approved resources rather than receiving broad administrative access.

Short-lived credentials or role-based authentication reduce the exposure associated with static credentials stored in pipeline configuration. The pipeline can obtain temporary authorization at runtime through an approved identity mechanism.

Using a shared administrator account prevents accountability and gives the pipeline excessive privilege. Disabling audit logging removes evidence needed to investigate changes and does not improve deployment security.

QUESTION NO: 6

Which of the following is a primary advantage of using an orchestration model like N/E/S/W-bound APIs?

Choose one:

- A. It makes coffee for the IT team
- B. It increases the number of office plants
- C. It simplifies and automates the creation and management of services
- D. It predicts the weather accurately

ANSWER: C

Explanation:

It simplifies and automates the creation and management of services is correct. An orchestration model coordinates infrastructure, network, and application components through programmatic interfaces, allowing a higher-level system to define a desired service and automate the steps required to deploy and maintain it. Northbound APIs commonly let portals, automation pipelines, and business-level systems request services, while southbound APIs enable the orchestrator or controller to configure the underlying platforms and devices. East-west interfaces can support communication and integration among peer systems or services. Together, these API-driven integrations reduce manual configuration work, improve repeatability, and make service delivery faster and more consistent across environments. In an F5 cloud context, orchestration and declarative automation can consistently provision application delivery and security services alongside the application infrastructure. F5 describes its automation tooling as enabling declarative configuration and repeatable deployments, while the broader cloud-native approach relies on APIs to automate operational workflows. See the [F5 Declarative Onboarding documentation](#) and the [F5 orchestration glossary](#) for additional context.

QUESTION NO: 7

Scenario: A high-availability pair of BIG-IP Virtual Edition instances is deployed across cloud availability zones. During a failover, the active unit must update cloud networking resources so that application traffic reaches the newly active instance.

Which F5 component is designed to automate this cloud resource failover process?

Choose one:

- A. F5 Declarative Onboarding
- B. F5 Cloud Failover Extension
- C. F5 Application Services 3 Extension
- D. BIG-IQ Centralized Management

ANSWER: B

Explanation:

F5 Cloud Failover Extension is designed to use cloud-provider APIs to update cloud networking objects during BIG-IP failover. Depending on the cloud deployment model, this can include moving or updating addresses, routes, forwarding rules, or other cloud resources required to direct traffic to the active BIG-IP instance.

AS3 manages application-service declarations, while Declarative Onboarding configures initial BIG-IP system settings. BIG-IQ can centrally manage BIG-IP devices but is not the focused cloud-resource failover mechanism in this scenario.

QUESTION NO: 8

In the N/E/S/W-bound API model, what does "N" typically represent?

Choose one:

- A. Noodle
- B. Narrow
- C. Network
- D. North

ANSWER: C D

Explanation:

In directional API terminology, "N" represents **North**, as in a northbound API or northbound interface. Directional labels describe the relationship between software layers or systems rather than a physical compass direction. A northbound interface generally exposes services, data, status, or control capabilities from a lower-layer platform to higher-level

consumers, such as orchestration, management, analytics, or business applications. In F5 cloud and application-delivery environments, APIs are commonly used to let automation and management tooling interact with application and infrastructure services; the direction indicates the role and relative position of the connected systems. “North” is therefore the conventional expansion of the N in the N/E/S/W-bound API model. The word “Network” can be associated with the domain in which these interfaces operate, but the directional abbreviation itself is North. For general F5 API documentation and API access concepts, see [F5 Distributed Cloud API documentation](#). For an overview of F5 automation and API-oriented management capabilities, see [F5 API documentation](#).