



Up-to-date Practice Test with Latest Questions and Answers covering latest syllabus and topics of the exam. Makes you ready to face actual exam.



C1000-176 Practice Questions  
C1000-176 Practice Test  
C1000-176 Practice Exam  
C1000-176 Exam Questions  
C1000-176 Study Guide



[killexams.com](https://killexams.com)

**IBM**

# C1000-176

IBM Certified Advanced Architect - Cloud v2 (C9006300)

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/C1000-176>



### Question: 592

A pharmaceutical company is deploying a research application on IBM Cloud, using IBM Cloud Kubernetes Service and IBM Cloud Block Storage. The application stores proprietary data and requires compliance with industry standards. Which actions ensure secure storage and access, per IBM Cloud security best practices?

- A. Configure encryption for Block Storage and implement RBAC for Kubernetes
- B. Disable encryption for Block Storage and use public endpoints for Kubernetes
- C. Use a third-party encryption tool and deploy Kubernetes in a public subnet
- D. Enable default encryption for Block Storage and use API keys for Kubernetes

Answer: A

Explanation: Configuring encryption for Block Storage and implementing RBAC for Kubernetes ensures secure storage and access, aligning with IBM Cloud security best practices.

### Question: 593

A retail application on IBM Cloud VPC requires immutable infrastructure for its database tier. The architect uses Terraform to provision IBM Cloud Databases for PostgreSQL. Which configuration ensures a new database instance is created with updated settings before the old one is deleted?

- A. 

```
resource "ibm_database" "postgres" { database = "postgresql" version = "13" lifecycle { create_before_destroy = true } }
```
- B. 

```
resource "ibm_database" "postgres" { database = "postgresql" version = "13" lifecycle { ignore_changes = [version] } }
```
- C. 

```
resource "ibm_database" "postgres" { database = "postgresql" version = "13" lifecycle { prevent_destroy = true } }
```
- D. 

```
resource "ibm_database" "postgres" { database = "postgresql" version = "13" lifecycle { replace_triggered_by = [version] } }
```

Answer: A

Explanation: Immutable infrastructure requires creating a new database instance before deleting the old one. The lifecycle { create\_before\_destroy = true } in the ibm\_database resource ensures this. Ignoring version changes prevents updates, preventing destruction stops deletion, and replace\_triggered\_by is not a standard lifecycle rule for this use case.

**Question: 594**

A logistics company is architecting a real-time tracking system on IBM Cloud that must auto-heal after a node failure in a Red Hat OpenShift cluster. The system relies on IBM Event Streams for messaging. Which configuration ensures auto-healing with minimal disruption?

- A. Configure IBM Event Streams with automated topic replication across zones
- B. Deploy the application with OpenShift's built-in replication controller and liveness probes
- C. Use IBM Cloud Activity Tracker to monitor node failures and trigger manual recovery
- D. Enable IBM Cloud Backup for OpenShift to restore failed nodes

Answer: B

Explanation: OpenShift's replication controller ensures that a specified number of pod replicas are running, automatically replacing failed pods on healthy nodes. Liveness probes detect unhealthy pods and trigger restarts, enabling auto-healing. IBM Event Streams' topic replication ensures messaging reliability but does not address node failure recovery. IBM Cloud Activity Tracker is for auditing, not recovery, and IBM Cloud Backup is for data restoration, not real-time node recovery.

**Question: 595**

A global law firm is architecting a hybrid multicloud platform for case management, requiring secure storage across IBM Cloud, Azure, and on-premises systems with data residency in the UK. Which components ensure compliance and security? (Select all that apply)

- A. Azure SQL Database for case storage
- B. IBM Cloud Hyper Protect DBaaS in the UK
- C. IBM Cloud Satellite to host services in the UK
- D. IBM Cloud Pak for Data for case analytics

Answer: B, C

Explanation: UK data residency requires secure, local storage. IBM Cloud Hyper Protect DBaaS in the UK provides encrypted storage, and IBM Cloud Satellite extends compliant services to UK data centers. Azure SQL Database may not ensure residency, and Pak for Data focuses on analytics.

**Question: 596**

A utility company's grid monitoring system on IBM Cloud must achieve 99.999% availability across zone outages. The system uses IBM Cloud Kubernetes Service (IKS) and IBM Cloud File Storage.

Which approach ensures this level of resilience?

- A. Configure IKS with worker nodes spread across multiple zones
- B. Enable cross-region replication for file storage
- C. Deploy IKS in a single zone with pod auto-scaling
- D. Use IBM Cloud Backup for file storage snapshots

Answer: A

Explanation: Achieving 99.999% availability requires distributing workloads across multiple zones in a multi-zone region. Configuring IKS with worker nodes spread across zones ensures that zone outages do not disrupt the cluster, and Kubernetes reschedules pods to healthy nodes. Cross-region replication and backups are for disaster recovery, not zone-level resilience. Single-zone IKS is insufficient for high availability.

**Question: 597**

A financial institution is designing a payment processing system on IBM Cloud that must auto-heal after a node failure in a Red Hat OpenShift cluster. The system uses IBM Cloud Event Streams. Which configuration ensures minimal disruption?

- A. Configure IBM Event Streams with automated topic replication
- B. Deploy the application with OpenShift's Horizontal Pod Autoscaling and readiness probes
- C. Use IBM Cloud Activity Tracker to monitor failures
- D. Enable IBM Cloud Backup for OpenShift

Answer: B

Explanation: OpenShift's Horizontal Pod Autoscaling replaces failed pods, and readiness probes ensure only healthy pods receive traffic, minimizing disruption. Event Streams replication, Activity Tracker, and backups do not address node failure recovery.

**Question: 598**

A manufacturing company is deploying IoT workloads on IBM Cloud, requiring automated scaling based on device data volume. Which solution ensures operational efficiency? (Select all that apply)

- A. IBM Cloud Code Engine for auto-scaling workloads
- B. IBM Cloud Functions for event-driven processing
- C. IBM Cloud Kubernetes Service with HPA
- D. IBM Watson IoT Platform for device management

Answer: A, C

Explanation: IBM Cloud Code Engine provides auto-scaling for containerized or serverless workloads,

ideal for IoT data processing. IBM Cloud Kubernetes Service with Horizontal Pod Autoscaling (HPA) scales containerized applications based on metrics like data volume. IBM Cloud Functions is event-driven but lacks comprehensive scaling for IoT workloads. IBM Watson IoT Platform manages devices but doesn't handle workload scaling.

**Question: 599**

An e-commerce platform is migrating a MongoDB database to IBM Cloud Databases for MongoDB Enterprise. The database requires sharding for scalability. Which command configures a sharded cluster?

- A. `mongod --shardsvr --replSet rs0 --port 27017`
- B. `mongos --configdb configReplSet/cfg1:27017,cfg2:27017 --port 27017`
- C. `mongo --eval "sh.enableSharding('myDB')"`
- D. `mongod --configsvr --replSet configReplSet --port 27017`

Answer: B

Explanation: The `mongos` command configures the query router for a sharded MongoDB cluster, specifying the config server replica set (`configReplSet`) and port. This is essential for sharding in IBM Cloud Databases for MongoDB Enterprise. Option A starts a shard server, D starts a config server, and C enables sharding on a database but does not configure the cluster.

**Question: 600**

A media company is migrating a digital asset management system to IBM Cloud. The system runs on VMs and requires global content delivery. The architect must design a solution for minimal disruption. Which IBM Cloud services should be used? (Select all that apply)

- A. IBM Cloud Internet Services for CDN
- B. IBM Cloud Migration Services for VM migration
- C. IBM Cloud Pak for Data for asset analytics
- D. IBM Cloud Virtual Servers for VSI

Answer: A, B, D

Explanation: IBM Cloud Internet Services provides CDN for global content delivery. IBM Cloud Migration Services ensure minimal-disruption VM migration. IBM Cloud Virtual Servers for VSI host the system's VMs. IBM Cloud Pak for Data is not relevant, as it focuses on analytics, not asset management migration.

**Question: 601**

A financial services firm is designing a hybrid multicloud platform with IBM Cloud VPC and AWS. The application requires secure file storage. Which two configurations provide encrypted storage?

- A. Configure AWS EFS with KMS encryption

- B. Deploy IBM Cloud File Storage with BYOK
- C. Use IBM Cloud Object Storage with default encryption
- D. Implement AWS S3 with SSE-S3

Answer: A,B

Explanation: AWS EFS with KMS encryption and IBM Cloud File Storage with Bring Your Own Key (BYOK) provide customer-controlled, encrypted file storage solutions. IBM Cloud Object Storage with default encryption and AWS S3 with SSE-S3 use provider-managed keys, which do not meet the requirement for customer-managed encryption.

### Question: 602

A logistics company is designing a supply chain visibility platform on IBM Cloud. The platform must integrate IoT data, support real-time tracking, and store data for analytics. Which architectural pattern and service combination is most suitable?

- A. Event-driven architecture with IBM Cloud Functions
- B. Microservices architecture with IBM Cloud Kubernetes Service
- C. Monolithic architecture with IBM Cloud Virtual Servers
- D. Serverless architecture with IBM Cloud Code Engine

Answer: B

Explanation: A microservices architecture with IBM Cloud Kubernetes Service supports real-time IoT data processing, scalability for tracking, and integration with analytics services. Event-driven or serverless architectures are less suitable for complex, stateful systems. A monolithic architecture lacks flexibility.

### Question: 603

A bank is modernizing a payment processing application to IBM Cloud Kubernetes Service. The application requires secure communication using mutual TLS (mTLS). Which Istio configuration enables mTLS?

- A. `kubectl apply -f <(istioctl kube-inject -f deployment.yaml)`
- B. `istioctl install --set profile=demo --set meshConfig.authPolicy=MUTUAL_TLS`
- C. `kubectl create -f mutual-tls-policy.yaml`
- D. `istioctl upgrade --set global.mtls.enabled=true`

Answer: B

Explanation: The `istioctl install` command with `meshConfig.authPolicy=MUTUAL_TLS` configures Istio to enforce mTLS across the service mesh, ensuring secure communication for the payment application. Option A injects Istio sidecars without mTLS configuration. Option C is vague and lacks specific mTLS details. Option D is for upgrades, not initial configuration.

## Istio mTLS Configuration Command

plain

Show inline

### Question: 604

A healthcare provider needs to migrate a patient management system to IBM Cloud while adhering to HIPAA compliance. The system includes a mix of VMs running on VMware and bare metal servers with custom configurations. The architect must design a solution to migrate this workload with minimal disruption. Which tools and services should be used? (Select all that apply)

- A. IBM Cloud Migration Services for VM migration
- B. IBM Cloud Bare Metal Servers for custom configurations
- C. IBM Cloud Pak for Data for data analytics
- D. VMware HCX for hybrid cloud mobility

Answer: A, B, D

Explanation: IBM Cloud Migration Services facilitate the migration of VMs to IBM Cloud, ensuring minimal disruption. IBM Cloud Bare Metal Servers support custom configurations required by the patient management system, maintaining performance and HIPAA compliance. VMware HCX enables seamless hybrid cloud mobility for VMware-based VMs, ensuring compatibility and minimal downtime. IBM Cloud Pak for Data is irrelevant for migration, as it focuses on data analytics, not workload migration.

### Question: 605

A banking institution is designing a fraud detection system on IBM Cloud that must auto-heal after a node failure in a Red Hat OpenShift cluster. The system uses IBM Watson for AI processing. Which configuration ensures minimal downtime?

- A. Configure IBM Watson with automated scaling
- B. Deploy the application with OpenShift's replication controller and liveness probes
- C. Use IBM Cloud Activity Tracker to monitor node failures
- D. Enable IBM Cloud Backup for OpenShift

Answer: B

Explanation: OpenShift's replication controller replaces failed pods, and liveness probes restart unhealthy ones, ensuring auto-healing with minimal downtime. Watson's scaling addresses AI workload but not node failures. Activity Tracker and backups are for auditing and recovery, not real-time healing.

### Question: 606

An architect is designing a centralized logging solution for a hybrid multicloud environment. The solution must aggregate logs from IBM Cloud Functions, Red Hat OpenShift, and an AWS EC2 instance, with a retention period of 180 days. Which configuration meets these requirements?

- A. Deploy IBM Log Analysis with LogDNA in IBM Cloud and configure Fluentd to forward logs from OpenShift and AWS EC2
- B. Enable IBM Cloud Activity Tracker to collect logs and integrate with IBM Event Streams for cross-cloud log aggregation
- C. Use IBM Cloud Monitoring with Sysdig to collect logs and store them in IBM Cloud Object Storage
- D. Configure AWS CloudWatch to forward logs to IBM Log Analysis via a Lambda function

Answer: A

Explanation: IBM Log Analysis with LogDNA supports log aggregation and retention (configurable up to 180 days). Fluentd is a versatile log forwarder that can collect logs from OpenShift and AWS EC2, integrating them with IBM Log Analysis. IBM Cloud Activity Tracker is for audit logs, not application logs. Sysdig is for metrics, not logs, and AWS CloudWatch integration via Lambda is not a standard approach for IBM Log Analysis.

### Question: 607

A retail chain is deploying a loyalty program application on IBM Cloud, requiring secure API access and compliance with GDPR. The application uses Kubernetes and public APIs. Which services should the architect implement?

- A. IBM Cloud App ID, IBM Cloud Key Protect, IBM Cloud Activity Tracker
- B. IBM Cloud Hyper Protect DBaaS, IBM Cloud Flow Logs, IBM Cloud Monitoring
- C. IBM Cloud Secrets Manager, IBM Cloud Transit Gateway, IBM Cloud Log Analysis
- D. IBM Cloud Certificate Manager, IBM Cloud Security Advisor, IBM Cloud IAM

Answer: A

Explanation: For GDPR compliance and secure APIs, IBM Cloud App ID provides OAuth-based authentication for Kubernetes APIs. IBM Cloud Key Protect encrypts customer data at rest, supporting GDPR. IBM Cloud Activity Tracker logs API access, ensuring auditability. The other options are less suitable: Hyper Protect DBaaS is database-specific; Flow Logs monitor traffic; and Certificate Manager is for TLS, not API security.

### Question: 608

A healthcare provider is migrating its patient records system to a hybrid multicloud environment using IBM Cloud and AWS. The system must support portable workloads and comply with GDPR. Which approach ensures portability?

- A. Deploy the system using IBM Cloud Pak for Applications on Red Hat OpenShift
- B. Implement AWS Fargate for containerized workloads
- C. Use IBM Cloud Hyper Protect DBaaS for data storage
- D. Rely on virtual machines with custom configurations

Answer: A

Explanation: For workload portability in a hybrid multicloud environment with GDPR compliance, the solution must provide a standardized, containerized platform with compliance features. IBM Cloud Pak for Applications, built on Red Hat OpenShift, offers a portable, Kubernetes-based environment that runs consistently across IBM Cloud and AWS, with security features to support GDPR compliance. AWS Fargate is specific to AWS, limiting portability. IBM Cloud Hyper Protect DBaaS focuses on secure databases, not workload portability. Virtual machines with custom configurations reduce portability.

**Question: 609**

A telecommunications provider is designing a hybrid multicloud network for its 5G core services, hosted on IBM Cloud VPCs and on-premises infrastructure. The network must support high availability and low-latency communication. Which design considerations should the architect implement? (Select all that apply)

- A. Configure IBM Cloud Transit Gateway with redundancy
- B. Deploy IBM Cloud Direct Link for private connectivity
- C. Implement a public internet-based load balancer
- D. Use IBM Cloud Flow Logs for real-time monitoring

Answer: A, B

Explanation: For a 5G core services network requiring high availability and low-latency communication, the design must prioritize private, redundant connectivity. IBM Cloud Transit Gateway with redundancy ensures high availability by interconnecting IBM Cloud VPCs and on-premises infrastructure with fault-tolerant routing. IBM Cloud Direct Link provides a private, low-latency connection between on-premises systems and IBM Cloud, critical for 5G performance. A public internet-based load balancer introduces latency and security risks. IBM Cloud Flow Logs are for monitoring, not directly supporting availability or latency.

**Question: 610**

A telecommunications company is migrating a billing application to IBM Cloud. The application, a Java monolith, interacts with a MongoDB database. The architect plans to containerize the application on IBM Cloud Red Hat OpenShift and migrate the database to IBM Cloud Databases for MongoDB. Which approach minimizes application downtime during database migration? (Select one)

- A. Export the MongoDB database and import it into the DBaaS using mongodump/mongorestore
- B. Set up MongoDB replica sets with the DBaaS as a secondary node and promote it after synchronization
- C. Use IBM Cloud Pak for Data to migrate the database schema and data
- D. Store database backups in IBM Cloud Object Storage and restore to the DBaaS

Answer: B

Explanation: Setting up MongoDB replica sets with IBM Cloud Databases for MongoDB as a secondary node allows real-time data synchronization. Once synchronized, promoting the DBaaS as the primary node minimizes downtime. Mongodump/mongorestore requires significant downtime for large datasets. Cloud Pak for Data is not suited for MongoDB migration. Using Object Storage for backups is slow and disruptive.



Killexams.com is a leading online platform specializing in high-quality certification exam preparation. Offering a robust suite of tools, including Exam Questions, practice tests, and advanced test engines, Killexams.com empowers candidates to excel in their certification exams. Discover the key features that make Killexams.com the go-to choice for exam success.



## Practice Exam Questions Based on Current Exam Objectives

Killexams.com provides practice exam questions aligned with the latest official exam objectives and latest syllabus. Our content is reviewed and updated regularly to reflect recent changes announced by certification vendors. By studying these practice questions, candidates will cover the structure, difficulty level, and topics of the actual exam, helping them prepare more effectively and efficiently.

## Comprehensive Practice Exams (PDF Format)

Killexams.com offers multiple-choice questions (MCQs) in easy-to-read PDF format, covering all major domains of the exam. Each PDF contains a structured collection of practice questions and verified answers designed to support focused study. These MCQs help candidates reinforce key concepts, identify knowledge gaps, and improve exam readiness through consistent practice.

## Realistic Practice Tests (Online Test Engine & Desktop Test Engine)

To support hands-on preparation, Killexams.com provides practice tests through both an Online Test Engine and a Desktop Test Engine. These tools are designed to simulate a real exam environment, allowing candidates to practice under exam-like conditions, with latest syllabus and topics of the exam. Performance tracking, test history, and result analysis help users evaluate their progress and focus on areas that need improvement.

## Risk-Free Purchase Policy

Killexams.com follows a transparent and customer-friendly purchase policy. If users are not satisfied with the study materials, they may request assistance or a refund in accordance with our published terms and conditions. This policy reflects our commitment to customer satisfaction, fairness, and confidence in our preparation resources.

## Regularly Updated Content

Our practice question bank is reviewed and updated on an ongoing basis to stay aligned with the latest exam outlines and vendor updates. This ensures candidates are studying up-to-date, relevant material, and preparing with content that reflects current exam expectations, helping them stay confident and well-prepared.