



Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt. ----- Guaranteed.



CLF-C02 Dumps
CLF-C02 Braindumps
CLF-C02 Real Questions
CLF-C02 Practice Test
CLF-C02 Actual Questions



killexams.com

Amazon

CLF-C02

AWS Certified Cloud Practitioner 2025

ORDER FULL VERSION

<https://killexams.com/pass4sure/exam-detail/CLF-C02>



Question: 870

A company with a Business Support plan needs help resolving an Amazon EKS cluster failure, requiring 24/7 support, cost optimization insights, and proactive health checks. Which AWS support service provides this?

- A. AWS Managed Services with CloudWatch
- B. AWS IQ with Enterprise Support
- C. AWS Support with Business Support plan
- D. AWS Activate with Trusted Advisor

Answer: C

Explanation: The Business Support plan offers 24/7 support, Trusted Advisor for cost optimization, and health checks via Personal Health Dashboard. IQ is for experts, AMS is operational, and Activate is for startups.

Question: 871

You need to list log files in 's3://my-bucket/logs/' with 'fail' in the name and calculate their total size in MB. Which CLI command is correct?

- A. `aws s3api get-objects --bucket my-bucket --path logs/ --filter fail --output sizeMB`
- B. `aws s3 ls s3://my-bucket/logs/*fail* | sum Size / 1024 / 1024`
- C. `aws s3api list-objects-v2 --bucket my-bucket --prefix logs/ --query 'Contents[?contains(Key, fail)].Size' | jq 'reduce .[] as $x (0; . + $x) / 1048576'`
- D. `aws s3 list s3://my-bucket/logs/ --match fail | total / 1048576`

Answer: C

Explanation: list-objects-v2 with --query and jq sums sizes and converts to MB (divide by 1048576). Option A is correct; others are invalid or less precise.

Question: 872

An organization is architecting a highly resilient application on AWS and needs to ensure that its infrastructure can withstand regional failures while maintaining compliance with data residency laws in multiple countries. The application must also serve users globally with minimal latency. Which AWS architectural approach best aligns with these requirements, considering the relationships among Regions, Availability Zones, and edge locations?

- A. Deploy the application in a single Region with multiple Availability Zones and use Amazon CloudFront for global content delivery
- B. Deploy in a single Availability Zone with edge locations to cache content and handle failover
- C. Use multiple Regions with synchronized data replication and leverage AWS Global Accelerator for latency optimization
- D. Use a single Region with no edge locations and rely on Availability Zone redundancy for resilience

Answer: C

Explanation: Multiple Regions allow for disaster recovery, business continuity, and compliance with data sovereignty laws by keeping data in specific geographic locations. AWS Global Accelerator optimizes latency by routing traffic through edge locations to the nearest Region, enhancing global performance. Option A limits resilience to a single Region, C restricts scalability and redundancy, and D lacks latency optimization and global reach.

Question: 873

Your team needs to deploy a machine learning model training job that processes 10 TB of data with GPU acceleration. The job requires 16 GPUs and 256 GiB of memory. Which Amazon EC2 instance type is best suited for this workload?

- A. M6i.32xlarge
- B. G5.12xlarge
- C. C6g.16xlarge
- D. P4d.24xlarge

Answer: D

Explanation: The P4d.24xlarge instance provides 8 NVIDIA A100 GPUs (scalable to 16 via clustering), 1.1 TiB of memory, and is optimized for ML training. G5 is less powerful, C6g lacks GPUs, and M6i is general-purpose, not ML-optimized.

Question: 874

A company is developing a customer support application that requires a conversational

interface to handle user inquiries about product availability and pricing. The solution must integrate natural language processing to interpret user intents and provide real-time responses, while also supporting integration with a backend inventory system. Which AWS AI/ML service should the company use to build this conversational interface with minimal development effort?

- A. Amazon Kendra with search indexing enabled
- B. Amazon SageMaker with a custom NLP model trained on SageMaker JumpStart
- C. Amazon Lex with intent recognition and Lambda integration
- D. Amazon Comprehend with entity recognition and sentiment analysis

Answer: C

Explanation: Amazon Lex is designed for building conversational interfaces, such as chatbots, with natural language understanding (NLU) capabilities to recognize user intents and utterances. It supports integration with AWS Lambda for backend processing (e.g., querying an inventory system), making it ideal for this use case with minimal development effort. Amazon Kendra is an intelligent search service, not a conversational tool. Amazon SageMaker requires building and training custom models, increasing complexity. Amazon Comprehend analyzes text but doesn't provide conversational capabilities.

Question: 875

An organization is deploying a machine learning workload on AWS and must choose an interaction method for provisioning resources. The team is debating between the AWS CLI for scripted automation, the AWS Management Console for ease of use, or AWS CloudFormation for templated deployments. The workload requires consistent setup across multiple teams and regions. Which method should they select?

- A. AWS Management Console
- B. AWS CloudFormation
- C. AWS CLI
- D. AWS SDKs

Answer: B

Explanation: AWS CloudFormation ensures consistent, templated deployments across multiple teams and regions using infrastructure as code, making it the best choice for this machine learning workload. The AWS Management Console is manual and not scalable

for this use case. The AWS CLI and SDKs offer automation but lack the structured repeatability of CloudFormation.

Question: 876

A startup runs a web application with unpredictable traffic, requiring 5 t3.medium instances for baseline load and up to 15 additional instances during marketing campaigns that occur twice monthly. The team wants to balance cost savings with availability and avoid long-term commitments. Which compute purchasing option should they choose?

- A. 5 Reserved Instances with 1-year term and 15 Spot Instances
- B. 5 On-Demand Instances and 15 Savings Plans
- C. 5 Dedicated Instances and 15 On-Demand Instances
- D. 5 Spot Instances with Capacity Reservations and 10 Dedicated Hosts

Answer: A

Explanation: Reserved Instances (1-year) save costs for the baseline, while Spot Instances handle bursts cheaply with availability risks mitigated by the small scale. Savings Plans require commitment, Dedicated Instances are expensive, and Dedicated Hosts are unnecessary.

Question: 877

A company uses AWS Launch Wizard to deploy a Microsoft SQL Server on EC2 with a custom domain join script and needs to ensure the instance joins the domain only after DNS resolution is confirmed. Which configuration should they use?

- A. Add a "UserData" script with "Resolve-DnsName -Name domain.local" and a "DependsOn" clause for the DNS resource in the CloudFormation output
- B. Use a "CreationPolicy" with a "ResourceSignal" to wait for domain join completion
- C. Configure a "WaitCondition" to pause deployment until DNS resolves, triggered by a Lambda function
- D. Set a "Condition" in Launch Wizard to check DNS availability before EC2 provisioning

Answer: A

Explanation: Launch Wizard generates CloudFormation, where a "UserData" script (e.g., PowerShell's "Resolve-DnsName") can verify DNS resolution, and "DependsOn" ensures sequencing. "CreationPolicy" is for resource signals, "WaitCondition" is overkill with Lambda, and "Condition" isn't a Launch Wizard feature.

Question: 878

A company uses AWS Storage Gateway to provide on-premises access to product catalogs stored in Amazon S3. The solution must cache frequently accessed files locally and support Windows-based clients via SMB. Which Storage Gateway configuration should they deploy?

- A. Volume Gateway with cached volumes
- B. File Gateway with SMB and S3 backend
- C. Tape Gateway with S3 Glacier
- D. File Gateway with NFS and S3 backend

Answer: B

Explanation: File Gateway with SMB and an S3 backend caches files locally and supports Windows clients via SMB, meeting the requirements. Volume Gateway (option A) is for block storage, not SMB. Tape Gateway (option C) is for archival, and NFS (option D) is for Linux, not Windows.

Question: 879

A parent company uses AWS Organizations with 5 member accounts and purchases 20 c5.large Reserved Instances with a 1-year term and All Upfront payment in the management account. Two member accounts run 10 c5.large instances each, while the management account runs none. How does the Reserved Instance discount behave across the organization?

- A. The discount applies to one member account only, chosen randomly
- B. The discount applies only to the management account unless shared explicitly
- C. The discount splits evenly, covering 10 instances total across accounts
- D. The discount applies to all 20 instances across the two member accounts automatically

Answer: D

Explanation: In AWS Organizations, RIs purchased in the management account with Regional scope automatically apply to eligible instances across member accounts, covering all 20 c5.large instances. Option B requires explicit sharing only for account-specific RIs, C is incorrect, and D is random nonsense.

Question: 880

Your 50 TB app uses X-Ray to trace a Lambda function, requiring a segment filter for requests taking >500 ms and a custom annotation "env=prod". Which X-Ray command creates this filter?

- A. `aws xray configure-group --group MyGroup --filter "duration>0.5,env=prod"`
- B. `aws xray set-filter --name MyGroup --expr "lambda > 500ms env=prod"`
- C. `aws xray create-group --group-name MyGroup --filter-expression "service(lambdaC. { duration > 0.5 } annotation.env = \"prod\"")"`
- D. `aws xray trace-filter --name MyGroup --condition "lambda>500ms" --annotation "env=prod"`

Answer: C

Explanation: X-Ray's create-group uses filter expressions for duration (>500 ms = 0.5s) and annotations (env=prod). Option A is correct; others are invalid.

Question: 881

A gaming company uses Amazon RDS for its leaderboard database and must comply with GDPR, which requires data encryption and access auditing. The compliance requirements differ from those of Amazon EC2, which they previously used. How does the encryption responsibility shift when moving from EC2 to RDS?

- A. Customer fully manages encryption for both services
- B. Customer manages encryption on EC2, while AWS handles it for RDS
- C. AWS manages encryption for EC2, while customer handles it for RDS
- D. Encryption is optional for RDS but mandatory for EC2

Answer: A

Explanation: For both EC2 and RDS, customers manage encryption (e.g., EBS/KMS for EC2, KMS for RDS) under the shared responsibility model, as it's security in the cloud. AWS provides tools, but implementation is customer-driven, making A, C, and D incorrect.

Question: 882

An IoT solution processes telemetry data from thousands of devices using AWS IoT Core, with edge processing handled by AWS IoT Greengrass. The system must route messages to Amazon Kinesis Data Streams based on a rule evaluating the payload temperature > 75°C, then trigger a Lambda function for analytics. Which IoT Core rule syntax achieves this?

- A. `SELECT * WHERE temperature > 75 FROM 'devices/*' INTO kinesis_stream('my_stream')`
- B. `SELECT temperature FROM 'devices/#' WHERE temperature > 75 THEN kinesis_stream('my_stream')`
- C. `FROM 'devices/+telemetry' WHERE temperature > 75 TO kinesis_stream('my_stream')`
- D. `SELECT * FROM 'devices/+telemetry' WHERE temperature > 75 INTO kinesis_stream('my_stream')`

Answer: D

Explanation: AWS IoT Core rules use SQL-like syntax: `SELECT * FROM 'devices/+telemetry' WHERE temperature > 75` correctly filters messages with temperature > 75°C from a topic with wildcards (+), routing them to Kinesis via `kinesis_stream('my_stream')`.

Question: 883

A bioinformatics company needs to run CPU-intensive genomic sequencing workloads on AWS, requiring instances optimized for high-performance computing and low-latency interconnects for cluster processing. Which EC2 instance type is most appropriate?

- A. Compute-optimized (C6i) instances
- B. Memory-optimized (R5) instances
- C. Storage-optimized (H1) instances

D. General-purpose (M6g) instances

Answer: A

Explanation: Compute-optimized (C6i) instances offer high CPU performance and low-latency interconnects, ideal for genomic sequencing. Memory-optimized (R5) prioritize RAM, storage-optimized (H1) focus on I/O, and general-purpose (M6g) lack the CPU specialization needed.

Question: 884

A retailer's on-premises costs include \$50,000 fixed for hardware and \$2,000–\$5,000 variable for utilities monthly. On AWS, they use 10 m5.large instances (\$0.096/hour) 24/7. What's the cost shift?

- A. Both costs balance out
- B. Variable costs decrease due to efficiency
- C. Fixed costs rise with instance reservations
- D. Fixed costs drop, variable costs align with usage

Answer: D

Explanation: AWS: $10 \times \$0.096 \times 720 = \$691.20/\text{month}$ (\$8,294/year), all variable. Fixed (\$50,000) disappears, variable (\$24,000–\$60,000) shifts to usage-based pricing.

Question: 885

Your application uses the AWS SDK for Python (Boto3) to upload files to an S3 bucket. You need to implement multipart uploads for files exceeding 5 GB with automatic retries on failure. Which Boto3 method and parameter should you configure?

- A. `s3.upload_fileobj` with `TransferConfig(max_attempts=3)`
- B. `s3.put_object` with `MultipartThreshold=5GB`
- C. `s3.upload_file` with `Config=RetryAttempts(5)`
- D. `s3.create_multipart_upload` with `RetryMode=adaptive`

Answer: A

Explanation: Boto3's `s3.upload_fileobj` with `TransferConfig(max_attempts=3)` supports multipart uploads for files over 5 GB with retries. `upload_file` lacks retry config, `put_object` is for single uploads, and `create_multipart_upload` is low-level without retry parameters.

Question: 886

A gaming application on AWS requires a database to cache player profiles and game states in memory, delivering sub-millisecond latency for frequent updates, while persisting critical data to a durable store. Which memory-based database is most suitable?

- A. Amazon DynamoDB with DAX
- B. Amazon RDS with MySQL
- C. Amazon ElastiCache with Memcached
- D. Amazon Aurora

Answer: C

Explanation: Amazon ElastiCache with Memcached provides in-memory caching with sub-millisecond latency, suitable for game states, with persistence handled separately. DynamoDB with DAX is NoSQL-specific, RDS and Aurora are disk-based relational databases.

Question: 887

A retail chain needs to send transactional SMS alerts to customers when orders ship, alongside push notifications to their mobile app when items are back in stock, with the ability to process millions of messages daily and integrate with Amazon DynamoDB for customer preferences. Which AWS service should they use for this messaging system?

- A. AWS AppSync
- B. Amazon SES
- C. Amazon SQS
- D. Amazon Pinpoint

Answer: D

Explanation: Amazon Pinpoint supports high-volume SMS and push notifications with DynamoDB integration for preferences. SES is email-focused, SQS is for queuing, and AppSync is for GraphQL data sync, not alerts.

Question: 888

Your company stores petabytes of archival data in an S3 bucket, and you need to ensure that objects are protected against accidental deletion with a retention period of 5 years. Which S3 feature and configuration should you enable to enforce this requirement?

- A. S3 Glacier Vault Lock with LockPeriod=5y and legal hold
- B. S3 Versioning with RetentionDays=1825 and MFA Delete enabled
- C. S3 Object Lock with {"Mode":"Governance","RetentionPeriod":1825} in the bucket policy
- D. S3 Lifecycle with {"Protect":true,"Days":1825} and deletion disabled

Answer: C

Explanation: S3 Object Lock with Governance mode and a retention period (5 years = 1825 days) prevents deletion until the period expires. Versioning with MFA Delete adds protection but lacks retention, Glacier Vault Lock is for Glacier, and Lifecycle doesn't enforce retention. Option A is correct for S3.

Question: 889

A software development company is building a CI/CD pipeline for a new application and must decide how to interact with AWS services to deploy resources. The team is split between using the AWS SDKs for custom automation in their preferred programming language, the AWS Management Console for quick manual deployments, or AWS CloudFormation for templated infrastructure management. The pipeline requires consistent deployments across development, testing, and production environments. Which method should the team adopt?

- A. AWS CloudFormation
- B. AWS SDKs
- C. AWS Management Console
- D. AWS CLI

Answer: A

Explanation: AWS CloudFormation is the best fit for a CI/CD pipeline requiring consistent deployments across multiple environments. It uses infrastructure as code to define resources in templates, ensuring repeatability and consistency. The AWS Management Console is manual and error-prone for this use case, while the AWS SDKs and CLI are better for ad-hoc scripting but lack the structured, templated approach of CloudFormation.





KILLEXAMS.COM

Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:



Actual Exam Questions: Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.

Exam Dumps: Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.

Practice Tests: Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success: Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.

Updated Content: Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.

Technical Support: Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.