CERT-PASS

AWS AWS Certified Data Engineer – Associate DEA-C01

Free Practice Questions Preview

Here are 35 sample questions to help you get started. Unlock the full exam to access all 1040+ questions with detailed explanations.

Question 1 : Content Domain 1: Data Ingestion and Transformation

An energy provider is redesigning a data pipeline. It wants an operations email list and a chat integration to receive a pipeline-failure alert. The design should avoid unnecessary custom infrastructure. Choose the most appropriate AWS-native option. What should the data engineer do?

A. Create an Athena view for the alert.

B. Publish the alert to an Amazon SNS topic with the required subscriptions.

C. Store the alert in DynamoDB TTL records only.

D. Use an S3 multipart upload as the notification mechanism.

Answer: B
Reasoning: SNS topics fan out notifications to subscribed endpoints. Avoid this mistake: An Athena view is a query abstraction and cannot deliver operational notifications.

Question 2 : Content Domain 2: Data Store Management

A logistics platform is implementing a new analytics workflow. It needs single-digit millisecond order lookups by order ID at very high scale. The access pattern is known and joins are not required. The engineering team wants clear failure handling and observability. Choose the most appropriate AWS-native option. What is the most appropriate design?

A. Use Amazon DynamoDB with a suitable partition key.

B. Use Amazon Redshift for every operational request.

C. Use Athena for every customer-facing lookup.

D. Use Amazon Aurora and perform relational queries for every lookup even though the key-value access pattern is fixed.

Answer: A
Reasoning: DynamoDB is designed for scalable low-latency key-value and document access patterns. Avoid this mistake: Redshift is optimized for analytics rather than millisecond operational request paths.

Question 3 : Content Domain 3: Data Operations and Support

A pharmaceutical distributor needs to improve a data engineering workload. It needs a SQL report with a seven-day rolling average of daily sales. The team wants a supportable solution with clear ownership. The team wants to avoid unnecessary components. Which approach most directly meets the requirement?

A. Use a GROUP BY clause that returns one row per day without a window frame.

B. Use a PIVOT operation without an ordered window.

C. Use a window function over an ordered date range.

D. Use ROW_NUMBER() without aggregating the seven-day range.

Answer: C
Best choice: A SQL window function is the standard approach for rolling calculations across ordered rows. Why the alternative fails: A GROUP BY clause can produce daily totals, but it does not calculate a seven-day rolling window without an appropriate window function.

Question 4 : Content Domain 4: Data Security and Governance

A public-sector agency is implementing a new analytics workflow. It needs a Lambda role that can read objects only from one S3 landing prefix and cannot delete them. No AWS managed policy matches the requirement. A recent audit found that the existing process is difficult to operate consistently. The requirement is straightforward and the team prefers the simplest managed approach. What is the most appropriate design?

A. Make the bucket public-read.

B. Create a custom IAM policy that grants only the required s3:GetObject access to the landing prefix.

C. Attach AdministratorAccess to the role.

D. Create an IAM user with a long-lived access key and share it with the function.

Answer: B
Best choice: A custom policy can scope the action and resource to the exact least-privilege requirement. Why the main distractor is wrong: AdministratorAccess grants unnecessary permissions and violates least privilege.

Question 5 : Content Domain 1: Data Ingestion and Transformation

A media-streaming company is preparing a production data platform change. It uses several Glue crawlers and Glue ETL jobs. Jobs should start after catalog discovery completes. The engineering team wants clear failure handling and observability. The requirement is straightforward and the team prefers the simplest managed approach. What should the data engineer do?

A. Use an AWS Glue workflow with triggers.

B. Use an S3 storage-class transition as the job dependency.

C. Use a Redshift sort key.

D. Use a DynamoDB TTL attribute.

Answer: A
Best choice: Glue workflows and triggers coordinate Glue crawlers and ETL jobs. Why the main distractor is wrong: A lifecycle transition changes storage class and does not orchestrate crawler and job dependencies.

Question 6 : Content Domain 2: Data Store Management

An e-commerce company is reviewing an existing data-processing design. It has a document-oriented application that uses MongoDB-compatible APIs. The engineering team wants clear failure handling and observability. The team wants to avoid unnecessary components. Which approach most directly meets the requirement?

A. Use Amazon DocumentDB (with MongoDB compatibility).

B. Use Amazon Data Firehose.

C. Use Amazon Neptune.

D. Use Amazon Redshift Spectrum.

Answer: A
Correct: DocumentDB supports document workloads with MongoDB compatibility. Why the main distractor is wrong: Neptune is designed for graph relationships rather than MongoDB-compatible document APIs.

Question 7 : Content Domain 3: Data Operations and Support

An insurance company is redesigning a data pipeline. It stores months of structured application logs in S3 and needs occasional SQL investigations without maintaining a search cluster. A recent audit found that the existing process is difficult to operate consistently. The team wants to avoid unnecessary components. What should the data engineer do?

A. Use a DynamoDB local secondary index on the S3 bucket.

B. Use AWS Secrets Manager.

C. Use Amazon SNS to run SQL.

D. Use Amazon Athena to query the S3 logs.

Answer: D
Why this works: Athena is well suited to occasional serverless SQL analysis of structured logs in S3. Avoid this mistake: SNS distributes notifications and does not query S3 files.

Question 8 : Content Domain 1: Data Ingestion and Transformation

A public-sector agency is implementing a new analytics workflow. It wants automated unit tests, build steps, and promotion of pipeline code after approved commits. Operational simplicity is important after deployment. Which option directly satisfies the requirement? Which solution should the data engineer recommend?

A. Use AWS Glue crawlers as the release pipeline.

B. Use a CI/CD workflow such as AWS CodeBuild with AWS CodePipeline and infrastructure-as-code deployments.

C. Use an S3 Lifecycle rule to promote source code.

D. Deploy production changes manually from an engineer laptop only.

Answer: B
Best choice: A CI/CD workflow automates testing and repeatable promotion of pipeline changes. Why the main distractor is wrong: Manual laptop deployments weaken repeatability, auditing, and environment consistency.

Question 9 : Content Domain 4: Data Security and Governance

A SaaS company is reviewing an existing data-processing design. It is comparing a self-managed database on EC2 with an AWS managed database service. The team wants to reduce patching, backup administration, and routine maintenance. The team wants a supportable solution with clear ownership. Choose the most appropriate AWS-native option. What should the data engineer do?

A. Choose EC2 automatically because self-managed services always require less operational work.

B. Select an appropriate managed database service when its capabilities meet the workload requirements.

C. Store the database only in Lambda ephemeral storage.

D. Use an SNS topic as the database engine.

Answer: B
Best choice: Managed services can reduce routine operational responsibilities when they fit the workload. Trap: A self-managed EC2 database generally leaves more patching and administration with the customer.

Question 10 : Content Domain 2: Data Store Management

A healthcare analytics provider is implementing a new analytics workflow. It receives evolving JSON records in S3. A new optional field appears while older records do not contain it. The company wants to reduce recurring manual maintenance. The requirement is straightforward and the team prefers the simplest managed approach. Which option is the best fit?

A. Apply a backward-compatible schema-evolution strategy that treats the new field as optional.

B. Disable the catalog so queries cannot see schema differences.

C. Delete the existing dataset immediately.

D. Reject all historical partitions because the optional field is absent.

Answer: A
Recommended: Compatible optional additions can be handled while preserving historical usability. Trap: Rejecting all older records for an optional field creates unnecessary breakage.

Question 11 : Content Domain 1: Data Ingestion and Transformation

A media-streaming company is reviewing an existing data-processing design. It has bursty ingestion events. Downstream processing must absorb spikes, retry transient failures, and isolate poison messages for investigation. The company expects the workload to grow during seasonal peaks. Choose the most appropriate AWS-native option. Which option is the best fit?

A. Store failed messages only in Lambda ephemeral storage.

B. Invoke every worker synchronously from the producer and disable retries.

C. Place work messages in Amazon SQS and configure an appropriate dead-letter queue.

D. Use an SNS topic without any queueing or failure-isolation design.

Answer: C
Why this works: SQS decouples producers from consumers, buffers spikes, and can move repeatedly failing messages to a DLQ. Why the alternative fails: Synchronous invocation tightly couples producer success to downstream capacity

and makes bursts harder to absorb.

Question 12 : Content Domain 3: Data Operations and Support

A SaaS company is redesigning a data pipeline. It needs managed dashboards for business users to explore curated analytical results. The engineering team wants clear failure handling and observability. The requirement is straightforward and the team prefers the simplest managed approach. What is the most appropriate design?

A. Use Amazon QuickSight.

B. Use Amazon OpenSearch Dashboards even though the workload is general business analytics rather than indexed search.

C. Use Amazon Athena as the dashboarding layer without a BI tool.

D. Use AWS Glue DataBrew as the end-user dashboard service.

Answer: A
Correct: QuickSight provides managed business-intelligence dashboards and visualizations. Why the main distractor is wrong: Athena can query data, but it does not replace a managed BI dashboard and visualization service.

Question 13 : Content Domain 1: Data Ingestion and Transformation

A cybersecurity vendor is reviewing an existing data-processing design. It needs a managed path that buffers JSON events and delivers them to Amazon S3. No custom stream consumer or replay workflow is required. The design should avoid unnecessary custom infrastructure. The requirement is straightforward and the team prefers the simplest managed approach. Which solution should the data engineer recommend?

A. Use Amazon Kinesis Data Streams and operate a custom consumer application.

B. Use Amazon Data Firehose to deliver the records to Amazon S3.

C. Use Amazon SQS FIFO as the long-term analytics landing zone.

D. Use Amazon MSK and operate Kafka brokers and consumers.

Answer: B
Recommended: Data Firehose is designed for managed, buffered delivery to supported destinations such as Amazon S3. Why the alternative fails: Kinesis Data Streams is appropriate when custom consumers, replay, or specialized processing are required, but it adds avoidable work here.

Question 14 : Content Domain 2: Data Store Management

A manufacturing group is preparing a production data platform change. It stores temporary session records in DynamoDB that should expire automatically after 24 hours. The team is replacing a manual operational step. The team wants to avoid unnecessary components. Which approach most directly meets the requirement?

A. Create a Redshift materialized view.

B. Enable S3 Versioning.

C. Enable DynamoDB Time to Live (TTL) on an expiration attribute.

D. Increase Kinesis retention.

Answer: C
Best choice: DynamoDB TTL automatically removes expired items based on an item attribute. Avoid this mistake: S3 Versioning applies to S3 objects and does not expire DynamoDB records.

Question 15 : Content Domain 4: Data Security and Governance

A retail company is redesigning a data pipeline. It needs to grant analysts read-only access to selected Redshift schemas. The team wants a supportable solution with clear ownership. The requirement is straightforward and the team prefers the simplest managed approach. Which approach most directly meets the requirement?

A. Give every analyst the database administrator credentials.

B. Create and grant appropriate Redshift database roles or privileges.

C. Use an S3 Lifecycle policy.

D. Make the cluster publicly accessible without authentication.

Answer: B
Reasoning: Database roles and grants provide controlled access to Redshift objects. Trap: Sharing administrator credentials removes accountability and grants excessive authority.

Question 16 : Content Domain 3: Data Operations and Support

An e-commerce company is preparing a production data platform change. It has a very large dataset and needs a fast initial quality assessment before running complete validation where required. The design should avoid unnecessary custom infrastructure. The team wants to avoid unnecessary components. Which solution should the data engineer recommend?

A. Analyze a representative sample first, then run full checks for critical controls.

B. Skip sampling and approve the dataset without any validation.

C. Inspect only the first contiguous rows and treat them as representative of every partition.

D. Select records only from the smallest partition even though the distribution is uneven.

Answer: A
Reasoning: Sampling can provide an efficient initial assessment while critical rules may still require complete validation. Trap: Sampling only the first contiguous rows can bias the assessment and miss issues in other partitions.

Question 17 : Content Domain 1: Data Ingestion and Transformation

A travel-booking platform is redesigning a data pipeline. It must start an ingestion workflow at 02:00 UTC every day even if no new object event occurs. The design should avoid unnecessary custom infrastructure. The requirement is straightforward and the team prefers the simplest managed approach. Which approach most directly meets the requirement?

A. Create an AWS Glue crawler with no schedule and run it manually.

B. Increase the Kinesis retention period.

C. Configure an S3 ObjectCreated event only.

D. Create an Amazon EventBridge Scheduler schedule for the workflow.

Answer: D
Best choice: EventBridge Scheduler is designed for time-based invocation of supported targets. Trap: S3 events are driven by object activity and do not guarantee a daily run when no object arrives.

Question 18 : Content Domain 2: Data Store Management

A cybersecurity vendor is redesigning a data pipeline. It stores historical Parquet data in S3 and recent curated tables in Redshift. Analysts need joins without loading all history into the cluster. The company wants to reduce recurring manual maintenance. The team wants to avoid unnecessary components. What should the data engineer do?

A. Use DynamoDB Streams.

B. Use Amazon Data Firehose delivery retries.

C. Use Amazon Redshift Spectrum to query the external S3 data.

D. Use an S3 Lifecycle transition.

Answer: C
Best choice: Redshift Spectrum lets Redshift SQL query external data stored in S3. Why the main distractor is wrong: DynamoDB Streams exposes DynamoDB item changes and does not query historical Parquet data in S3.

Question 19 : Content Domain 1: Data Ingestion and Transformation

A retail company is preparing a production data platform change. It has a serverless workflow with Lambda validation, a Glue job, conditional branches, retries, and a final notification. The engineering team wants clear failure handling and observability. The team wants to avoid unnecessary components. Which approach most directly meets the requirement?

A. Use an Athena view as the workflow engine.

B. Use Amazon S3 Lifecycle rules.

C. Use one SNS topic as the complete orchestration state machine.

D. Use AWS Step Functions to orchestrate the workflow.

Answer: D
Recommended: Step Functions models sequencing, branching, retries, and service integrations for serverless workflows. Why the main distractor is wrong: SNS distributes messages but does not replace a workflow state machine with conditional execution and retries.

Question 20 : Content Domain 4: Data Security and Governance

An insurance company is redesigning a data pipeline. It must permit a Glue connection to reach an RDS database on the database port while blocking unnecessary inbound access. The workload will be used in a production environment. Choose the most appropriate AWS-native option. Which approach most directly meets the requirement?

A. Disable database encryption.

B. Make the database publicly accessible from every IPv4 address.

C. Add an S3 Lifecycle rule.

D. Update the relevant VPC security group rules to allow only the required traffic.

Answer: D
Correct: Security groups control network access and should permit only the required path and port. Why the alternative fails: Opening the database to all addresses creates unnecessary exposure.

Question 21 : Content Domain 3: Data Operations and Support

A media-streaming company is reviewing an existing data-processing design. It must summarize revenue by month and product category and show categories as report columns. A recent audit found that the existing process is difficult to operate consistently. The requirement is straightforward and the team prefers the simplest managed approach. What is the most appropriate design?

A. Use only an ORDER BY clause without aggregation.

B. Use grouping and pivoting in the analytical query or reporting layer.

C. Use a rolling window average without reshaping categories into columns.

D. Use a Redshift sort key as the report transformation.

Answer: B
Best choice: Grouping aggregates records, and pivoting reshapes categories into report columns where needed. Why the main distractor is wrong: An ORDER BY clause can sort results, but it does not aggregate or pivot the categories into the requested report shape.

Question 22 : Content Domain 2: Data Store Management

An insurance company is redesigning a data pipeline. It needs relational transactions, SQL compatibility, and managed high availability for an application database. The company expects the workload to grow during seasonal peaks. Which option directly satisfies the requirement? Which approach most directly meets the requirement?

A. Use Amazon Aurora.

B. Use Amazon OpenSearch Service as the system of record.

C. Use Amazon S3 Select as the transactional database.

D. Use Amazon Kinesis Data Streams as the relational store.

Answer: A
Best choice: Aurora is a managed relational database suitable for transactional SQL workloads. Avoid this mistake: S3 is object storage and does not provide relational transaction semantics.

Question 23 : Content Domain 1: Data Ingestion and Transformation

A travel-booking platform is preparing a production data platform change. It needs ongoing replication of inserts and updates from an Amazon Aurora MySQL database into an analytics landing zone with minimal source impact. A recent audit found that the existing process is difficult to operate consistently. The requirement is straightforward and the team prefers the simplest managed approach. What is the most appropriate design?

A. Use AWS Database Migration Service (AWS DMS) with change data capture (CDC).

B. Use Amazon AppFlow to capture database transaction-log changes.

C. Use an AWS Glue crawler as the change-capture mechanism.

D. Run a complete database export every few minutes.

Answer: A
Why this works: AWS DMS CDC is designed to replicate ongoing source changes after an initial load. Avoid this mistake: Repeated full exports increase source load and latency, and they are inefficient for incremental replication.

Question 24 : Content Domain 4: Data Security and Governance

An energy provider is redesigning a data pipeline. It has an application that connects to RDS and must store and rotate database credentials securely. A recent audit found that the existing process is difficult to operate consistently. The requirement is straightforward and the team prefers the simplest managed approach. What is the most appropriate design?

A. Store the password in a public S3 object.

B. Hard-code the password in source control.

C. Publish the credential to an SNS topic.

D. Use AWS Secrets Manager and configure rotation where supported.

Answer: D
Best choice: Secrets Manager is designed to store and rotate secrets securely. Why the main distractor is wrong: Hard-coded credentials are difficult to rotate and are vulnerable to accidental disclosure.

Question 25 : Content Domain 3: Data Operations and Support

An online marketplace is reviewing an existing data-processing design. It needs a reconciliation Lambda function to run every 15 minutes. Operational simplicity is important after deployment. Which option directly satisfies the requirement? Which solution should the data engineer recommend?

A. Use a Redshift materialized view as the invocation service.

B. Use S3 Object Lock retention periods.

C. Use Amazon Macie classification jobs as the scheduler.

D. Create an Amazon EventBridge Scheduler schedule.

Answer: D
Why this works: EventBridge Scheduler invokes supported targets on a recurring schedule. Avoid this mistake: Macie discovers sensitive data in S3 and is unrelated to workflow scheduling.

Question 26 : Content Domain 1: Data Ingestion and Transformation

A university research team is reviewing an existing data-processing design. It must pull data from a partner endpoint that accepts connections only from approved public IP addresses. The pipeline runs from a VP C. Operational simplicity is important after deployment. The requirement is straightforward and the team prefers the simplest managed approach. What should the data engineer do?

A. Route outbound traffic through a NAT gateway with an Elastic IP address and provide that address for the allowlist.

B. Use random public IP addresses from Lambda-managed networking.

C. Use an S3 gateway endpoint as the public source IP for the partner API.

D. Add a DynamoDB global secondary index.

Answer: A
Correct: A NAT gateway with an Elastic IP provides a stable public egress address that a partner can allowlist. Avoid this mistake: A gateway endpoint is for private connectivity to supported AWS services and does not become the public egress IP for a partner endpoint.

Question 27 : Content Domain 2: Data Store Management

A sports analytics company is preparing a production data platform change. It needs durable, scalable, cost-effective storage for raw sensor readings in multiple formats that several analytics services will query. A recent audit found that the existing process is difficult to operate consistently. Which option directly satisfies the requirement? Which option is the best fit?

A. Store all raw files directly in Amazon Redshift even when formats and access patterns are still evolving.

B. Store the raw data in Amazon S3.

C. Store the raw records in DynamoDB even though analysts need broad object-based lake access.

D. Store the raw files on an Amazon EFS file system as the primary analytics lake.

Answer: B
Correct: Amazon S3 is the common durable object-storage foundation for a data lake. Why the alternative fails: Loading all evolving raw data directly into Redshift is less flexible than using S3 as the durable lake landing zone.

Question 28 : Content Domain 1: Data Ingestion and Transformation

A manufacturing group is reviewing an existing data-processing design. It uploads objects to Amazon S3 throughout the day. A Lambda validator should run soon after each upload. The team is replacing a manual operational step. Choose the most appropriate AWS-native option. What is the most appropriate design?

A. Create an hourly EventBridge schedule that lists the whole bucket.

B. Run an AWS Glue crawler once per week.

C. Configure an S3 Lifecycle rule to transition the uploaded objects.

D. Configure an Amazon S3 ObjectCreated event notification to invoke the Lambda function.

Answer: D
Best choice: S3 event notifications can invoke Lambda in response to new objects. Trap: Scheduled bucket scans can

work but add delay and unnecessary list operations for this event-driven requirement.

Question 29 : Content Domain 2: Data Store Management

A food-delivery platform is implementing a new analytics workflow. It needs efficient bulk movement between S3 and Redshift for load and export operations. The company expects the workload to grow during seasonal peaks. The team wants to avoid unnecessary components. Which solution should the data engineer recommend?

A. Use DynamoDB Scan and PutItem calls.

B. Use row-by-row INSERT statements for every file and DELETE statements for export.

C. Use an SQS queue as the export format.

D. Use Redshift COPY for loading and UNLOAD for exporting.

Answer: D
Why this works: COPY and UNLOAD are purpose-built Redshift operations for bulk movement between S3 and Redshift. Why the alternative fails: Row-by-row inserts are inefficient for bulk loading, and DELETE does not export query results.

Question 30 : Content Domain 3: Data Operations and Support

An insurance company is reviewing an existing data-processing design. It wants centralized application logs from Lambda transformation functions. The architecture should remain understandable to the operations team. The team wants to avoid unnecessary components. Which approach most directly meets the requirement?

A. Use Amazon CloudWatch Logs.

B. Use DynamoDB TTL records only.

C. Use an S3 Lifecycle rule without a log-delivery path.

D. Use Amazon SNS message attributes as the durable log store.

Answer: A
Best choice: CloudWatch Logs is the standard centralized logging destination for Lambda application output. Avoid this mistake: SNS is a messaging service and not a durable searchable application-log store.

Question 31 : Content Domain 4: Data Security and Governance

A public-sector agency is reviewing an existing data-processing design. It wants another analytics team to query selected Redshift data without copying the full dataset. The workload will be used in a production environment. The requirement is straightforward and the team prefers the simplest managed approach. Which option is the best fit?

A. Copy every table manually each night regardless of need.

B. Use Redshift data sharing with appropriate permissions.

C. Share the cluster administrator password.

D. Export the entire dataset to a public bucket.

Answer: B
Best choice: Redshift data sharing enables governed access without duplicating the full dataset. Trap: Sharing administrator credentials violates least privilege and reduces accountability.

Question 32 : Content Domain 1: Data Ingestion and Transformation

An insurance company is reviewing an existing data-processing design. It needs a low-code scheduled transfer from Salesforce to Amazon S3. The workload will be used in a production environment. Which option directly satisfies the requirement? What is the most appropriate design?

A. Use Amazon AppFlow.

B. Use Amazon Kinesis Data Streams.

C. Use AWS DataSync.

D. Use AWS Database Migration Service (AWS DMS).

Answer: A
Correct: Amazon AppFlow provides managed integrations between supported SaaS applications and AWS services. Avoid this mistake: DataSync is for supported storage transfer paths, not Salesforce SaaS extraction.

Question 33 : Content Domain 2: Data Store Management

A public-sector agency is redesigning a data pipeline. It must keep raw objects frequently accessible for 30 days and retain them more cheaply for seven years. The team is replacing a manual operational step. Choose the most appropriate AWS-native option. Which approach most directly meets the requirement?

A. Move every object manually each month.

B. Configure an S3 Lifecycle transition to an appropriate storage class.

C. Use DynamoDB TTL on the S3 bucket.

D. Change the Redshift distribution style.

Answer: B
Best choice: S3 Lifecycle rules automate storage-class transitions based on object age. Why the alternative fails: Manual movement is error-prone and does not provide reliable lifecycle automation.

Question 34 : Content Domain 1: Data Ingestion and Transformation

A gaming studio needs to improve a data engineering workload. It must ingest fraud signals continuously. Consumers occasionally need to replay records after a processing defect is fixed. The architecture should remain understandable to the operations team. Which option directly satisfies the requirement? Which approach most directly meets the requirement?

A. Publish the events to an Amazon SNS topic without a durable stream.

B. Write directly to a Lambda function and discard the original events after invocation.

C. Use Amazon Kinesis Data Streams and configure a retention period that covers the replay window.

D. Use Amazon Data Firehose and rely on its destination retry duration as the replay mechanism.

Answer: C
Why this works: Kinesis Data Streams retains stream records so consumers can read them again during the configured retention window. Avoid this mistake: Data Firehose is a managed delivery service, but its retry behavior is not a general-purpose consumer replay log.

Question 35 : Content Domain 3: Data Operations and Support

A media-streaming company is redesigning a data pipeline. It needs to search recent Lambda logs for a specific exception and aggregate counts by function. The team wants a supportable solution with clear ownership. Which option directly satisfies the requirement? Which option is the best fit?

A. Use DynamoDB TTL.

B. Use an S3 Lifecycle transition.

C. Use CloudWatch Logs Insights.

D. Use AWS DMS CDC.

Answer: C
Recommended: CloudWatch Logs Insights provides interactive search and aggregation over CloudWatch Logs. Avoid this mistake: DMS CDC replicates database changes and does not query Lambda application logs.