- **Vendor:** Microsoft
- **Exam Code:** DP-200
- **Exam Name:** Implementing an Azure Data Solution
- **New Updated Questions from Braindump2go (Updated in May/2021)**

QUESTION 257
You have an Azure Data Factory that contains 10 pipelines.
You need to label each pipeline with its main purpose of either ingest, transform, or load. The labels must be available for grouping and filtering when using the monitoring experience in Data Factory.
What should you add to each pipeline?

A. a resource tag
B. a user property
C. an annotation
D. a run group ID
E. a correlation ID

Answer: C
Explanation:
Annotations are additional, informative tags that you can add to specific factory resources: pipelines, datasets, linked services, and triggers. By adding annotations, you can easily filter and search for specific factory resources.
Reference:
https://www.cathrinewilhelmsen.net/annotations-user-properties-azure-data-factory/

QUESTION 258
You have a data warehouse in Azure Synapse Analytics.
You need to ensure that the data in the data warehouse is encrypted at rest.
What should you enable?

A. Transparent Data Encryption (TDE)
B. Secure transfer required
C. Always Encrypted for all columns
D. Advanced Data Security for this database

Answer: A
Explanation:
Azure SQL Database currently supports encryption at rest for Microsoft-managed service side and client- side encryption scenarios.
Support for server encryption is currently provided through the SQL feature called Transparent Data Encryption.
Client-side encryption of Azure SQL Database data is supported through the Always Encrypted feature.
Reference:
https://docs.microsoft.com/en-us/azure/security/fundamentals/encryption-atrest

QUESTION 259
You have an Azure Blob storage account.
The storage account has an alert that is configured to indicate when the Availability metric falls below 100 percent.
You receive an alert for the Availability metric. The logs for the storage account show that requests are failing because of a ServerTimeoutError error.
What does ServerTimeoutError indicate?

A. Read and write storage requests exceeded capacity.
B. A transient server timeout occurred while the service was moved to a different partition to load balance requests.
C. A client application attempted to perform an operation and did not have valid credentials.
D. There was excessive network latency between a client application and the storage account.

Answer: D

QUESTION 260
You are monitoring an Azure Stream Analytics job by using metrics in Azure.
You discover that during the last 12 hours, the average watermark delay is consistently greater than the configured late arrival tolerance.
What is a possible cause of this behavior?

A. The job lacks the resources to process the volume of incoming data.
B. The late arrival policy causes events to be dropped.
C. Events whose application timestamp is earlier than their arrival time by more than five minutes arrive as inputs.
D. There are errors in the input data.

Answer: A
Explanation:
https://azure.microsoft.com/en-us/blog/new-metric-in-azure-stream-analytics-tracks-latency-of-your-streaming-pipeline/

QUESTION 261
You have an Azure Blob storage account.
Developers report that an HTTP 403 (Forbidden) error is generated when a client application attempts to access the storage account. You cannot see the error messages in Azure Monitor.
What is a possible cause of the error?

A. The client application is using an expired shared access signature (SAS) when it sends a storage request.
B. The client application deleted, and then immediately recreated a blob container that has the same name.
C. The client application attempted to use a shared access signature (SAS) that did not have the necessary permissions.
D. The client application attempted to use a blob that does not exist in the storage service.

Answer: C
Explanation:
https://docs.microsoft.com/en-us/rest/api/storageservices/sas-error-codes

QUESTION 262
You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1. Df1 contains a linked service.
You have an Azure Key vault named vault1 that contains an encryption key named key1.
You need to encrypt df1 by using key1.
What should you do first?

A. Disable purge protection on vault1.
B. Create a self-hosted integration runtime.
C. Disable soft delete on vault1.
D. Remove the linked service from df1.

Answer: D

Explanation:
Linked services are much like connection strings, which define the connection information needed for Data Factory to connect to external resources.
Incorrect Answers:
A, C: Data Factory requires two properties to be set on the Key Vault, Soft Delete and Do Not Purge
B: A self-hosted integration runtime copies data between an on-premises store and cloud storage.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key https://docs.microsoft.com/en-us/azure/data-factory/concepts-linked-services https://docs.microsoft.com/en-us/azure/data-factory/create-self-hosted-integration-runtime

QUESTION 263
You have an Azure subscription the contains the resources shown in the following table:

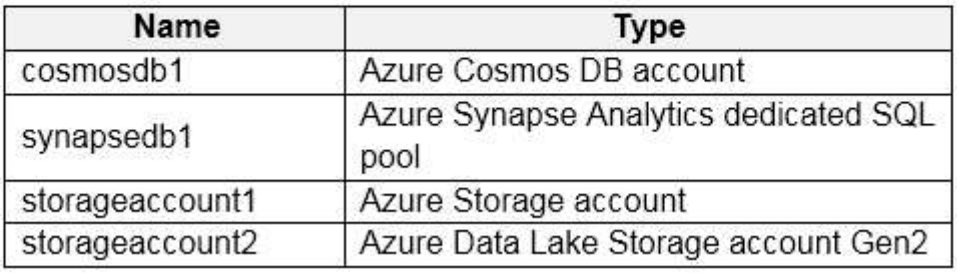

Name	Type
cosmosdb1	Azure Cosmos DB account
synapsedb1	Azure Synapse Analytics dedicated SQL pool
storageaccount1	Azure Storage account
storageaccount2	Azure Data Lake Storage account Gen2

All the resources have the default encryption settings.
You need to ensure that all the data stored in the resources is encrypted at rest.
What should you do?

A. Enable Azure Storage encryption for storageaccount1.
B. Enable Transparent Data Encryption (TDE) for synapsedb1.
C. Enable Azure Storage encryption for storageaccount2.
D. Enable encryption at rest for cosmosdb1.

Answer: B
Explanation:
Incorrect answers:
A, C: Azure Disks, and data in Azure Storage accounts are automatically encrypted at rest by default
D: All user data stored in Azure Cosmos DB is encrypted at rest by default
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/database/transparent-data-encryption-byok-overview
https://docs.microsoft.com/en-us/azure/synapse-analytics/security/workspaces-encryption
https://docs.microsoft.com/en-us/azure/cosmos-db/database-encryption-at-rest

QUESTION 264
You have an activity in an Azure Data Factory pipeline. The activity calls a stored procedure in a data warehouse in Azure Synapse Analytics and runs daily.
You need to verify the duration of the activity when it ran last.
What should you use?

A. the sys.dm_pdw_wait_stats data management view in Azure Synapse Analytics
B. an Azure Resource Manager template
C. activity runs in Azure Monitor
D. Activity log in Azure Synapse Analytics

Answer: C
Explanation:
Monitor activity runs. To get a detailed view of the individual activity runs of a specific pipeline run, click on the pipeline name.
Example:

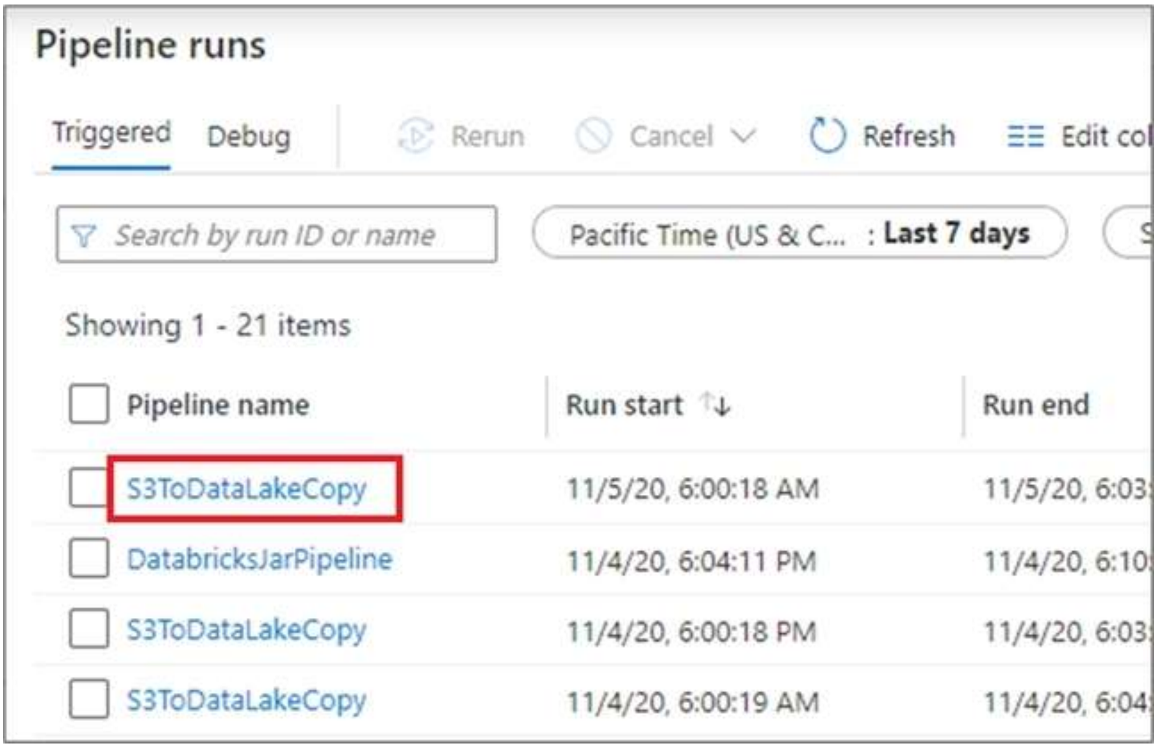

Pipeline runs

Search by run ID or name

Pacific Time (US & C... : Last 7 days

Showing 1 - 21 items

□	Pipeline name	Run start ↑↓	Run end
□	S3ToDataLakeCopy	11/5/20, 6:00:18 AM	11/5/20, 6:03
□	DatabricksJarPipeline	11/4/20, 6:04:11 PM	11/4/20, 6:10
□	S3ToDataLakeCopy	11/4/20, 6:00:18 PM	11/4/20, 6:03
□	S3ToDataLakeCopy	11/4/20, 6:00:19 AM	11/4/20, 6:04

The list view shows activity runs that correspond to each pipeline run. Hover over the specific activity run to get run-specific information such as the JSON input, JSON output, and detailed activity-specific monitoring experiences.

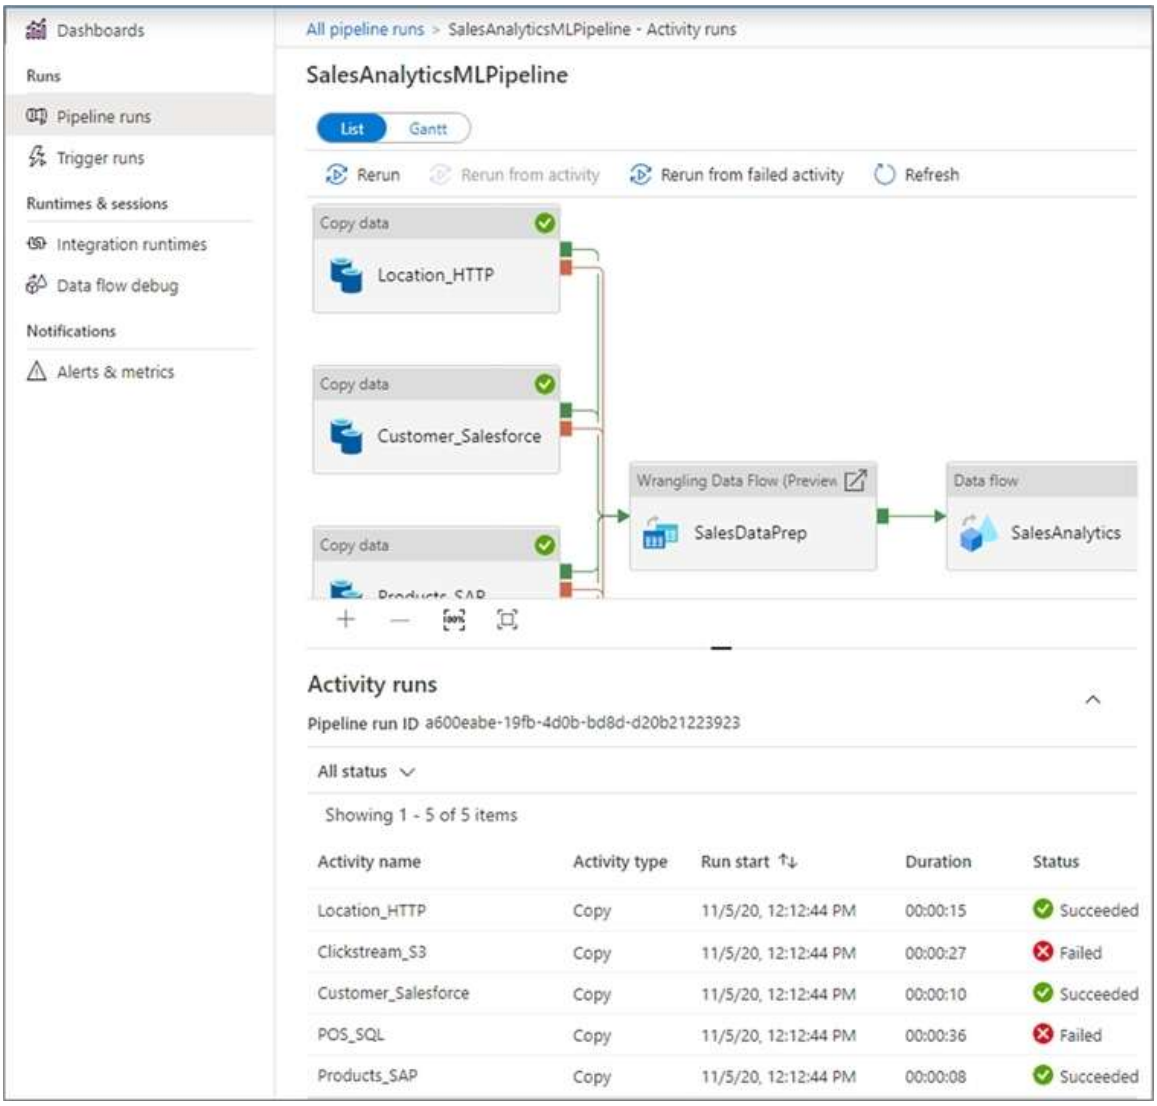

Activity runs

Pipeline run ID a600eabe-19fb-4d0b-bd8d-d20b21223923

Showing 1 - 5 of 5 items

Activity name	Activity type	Run start ⇅	Duration	Status
Location_HTTP	Copy	11/5/20, 12:12:44 PM	00:00:15	✓ Succeeded
Clickstream_S3	Copy	11/5/20, 12:12:44 PM	00:00:27	✗ Failed
Customer_Salesforce	Copy	11/5/20, 12:12:44 PM	00:00:10	✓ Succeeded
POS_SQL	Copy	11/5/20, 12:12:44 PM	00:00:36	✗ Failed
Products_SAP	Copy	11/5/20, 12:12:44 PM	00:00:08	✓ Succeeded

You can check the Duration.
Incorrect Answers:
A: sys.dm_pdw_wait_stats holds information related to the SQL Server OS state related to instances running on the different nodes.
Reference:
https://docs.microsoft.com/en-us/azure/data-factory/monitor-visually

QUESTION 265
You are monitoring an Azure Stream Analytics job.
The Backlogged Input Events count has been 20 for the last hour.
You need to reduce the Backlogged Input Events count.
What should you do?

A. Add an Azure Storage account to the job
B. Increase the streaming units for the job
C. Stop the job
D. Drop late arriving events from the job

Answer: B
Explanation:
General symptoms of the job hitting system resource limits include:
If the backlog event metric keeps increasing, it's an indicator that the system resource is constrained (either because of

output sink throttling, or high CPU).
Note: Backlogged Input Events: Number of input events that are backlogged. A non-zero value for this metric implies that your job isn't able to keep up with the number of incoming events. If this value is slowly increasing or consistently non-zero, you should scale out your job: adjust Streaming Units.
Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-scale-jobs https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-monitoring

QUESTION 266
Drag and Drop Question
You need to create an Azure Cosmos DB account that will use encryption keys managed by your organization.
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

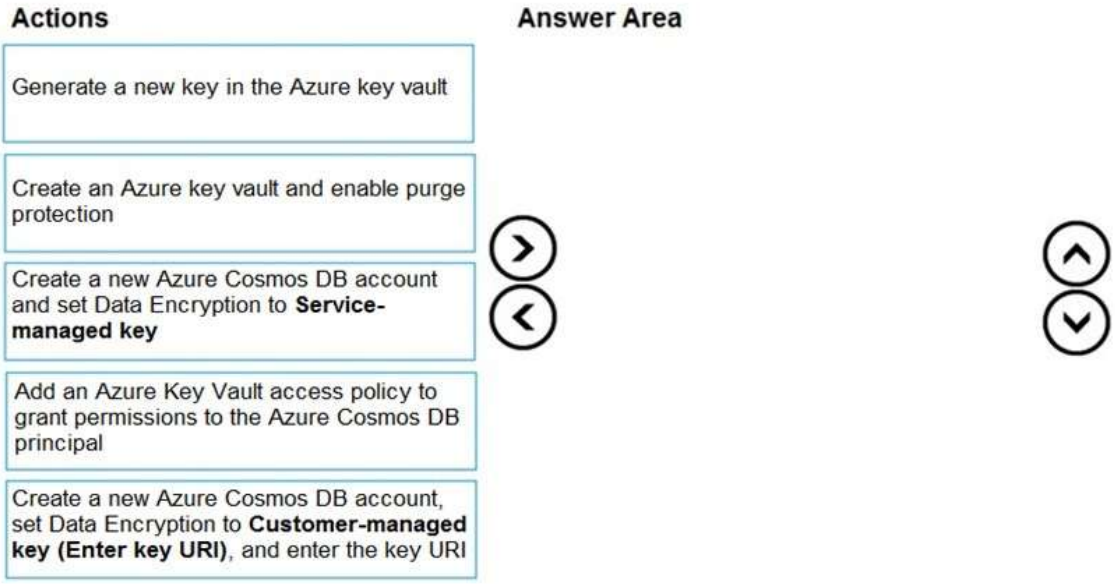

Actions

Generate a new key in the Azure key vault

Create an Azure key vault and enable purge protection

Create a new Azure Cosmos DB account and set Data Encryption to **Service-managed key**

Add an Azure Key Vault access policy to grant permissions to the Azure Cosmos DB principal

Create a new Azure Cosmos DB account, set Data Encryption to **Customer-managed key (Enter key URI)**, and enter the key URI

Answer Area

Answer:

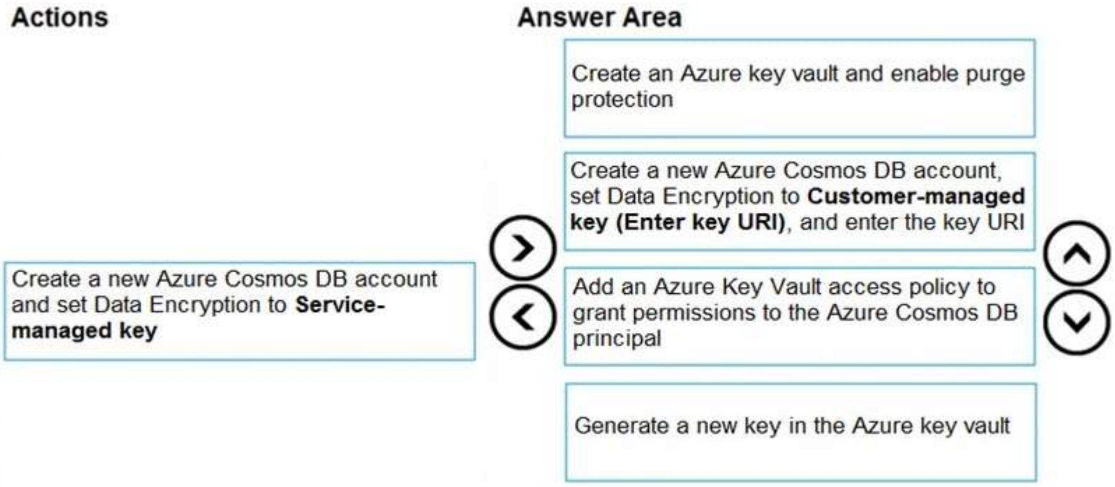

Actions

Create a new Azure Cosmos DB account and set Data Encryption to **Service-managed key**

Answer Area

Create an Azure key vault and enable purge protection

Create a new Azure Cosmos DB account, set Data Encryption to **Customer-managed key (Enter key URI)**, and enter the key URI

Add an Azure Key Vault access policy to grant permissions to the Azure Cosmos DB principal

Generate a new key in the Azure key vault

Explanation:
Step 1: Create an Azure key vault and enable purge protection Using customer-managed keys with Azure Cosmos DB requires you to set two properties on the Azure Key Vault instance that you plan to use to host your encryption keys:

Soft Delete and Purge Protection.
Step 2: Create a new Azure Cosmos DB account, set Data Encryption to Customer-managed Key (Enter key URI), and enter the key URI
Data stored in your Azure Cosmos account is automatically and seamlessly encrypted with keys managed by Microsoft (service-managed keys). Optionally, you can choose to add a second layer of encryption with keys you manage (customer-managed keys).
Step 3: Add an Azure Key Vault access policy to grant permissions to the Azure Cosmos DB principal
Add an access policy to your Azure Key Vault instance
Step 4: Generate a new key in the Azure key vault
Generate a key in Azure Key Vault
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/how-to-setup-cmk

QUESTION 267
Hotspot Question
You are building an Azure Stream Analytics query that will receive input data from Azure IoT Hub and write the results to Azure Blob storage.
You need to calculate the difference in readings per sensor per hour.
How should you complete the query? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

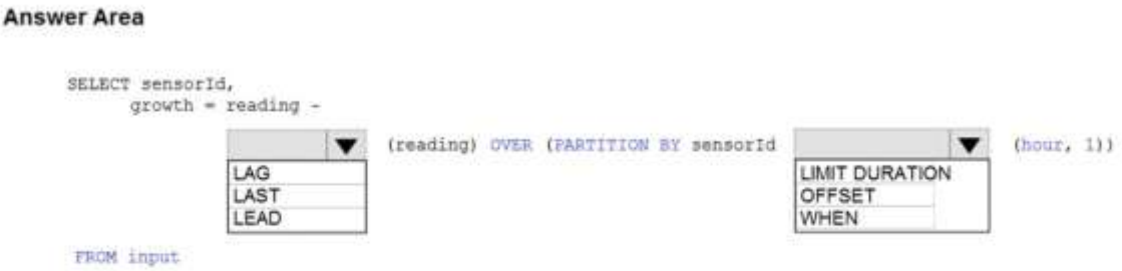
Answer Area

```
SELECT sensorId,
       growth = reading -
              [▼] (reading) OVER (PARTITION BY sensorId [▼] (hour, 1))
  FROM input
```

Answer:

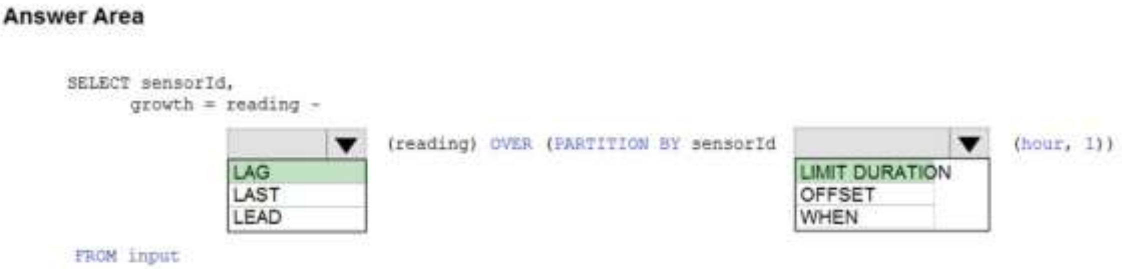
Answer Area

```
SELECT sensorId,
       growth = reading -
              [▼] (reading) OVER (PARTITION BY sensorId [▼] (hour, 1))
  FROM input
```

Explanation:
Box 1: LAG
The LAG analytic operator allows one to look up a "previous" event in an event stream, within certain constraints. It is very useful for computing the rate of growth of a variable, detecting when a variable crosses a threshold, or when a condition starts or stops being true.
Box 2: LIMIT DURATION
Example: Compute the rate of growth, per sensor:
SELECT sensorId,
growth = reading -
LAG(reading) OVER (PARTITION BY sensorId LIMIT DURATION(hour, 1)) FROM input
Reference:
https://docs.microsoft.com/en-us/stream-analytics-query/lag-azure-stream-analytics

QUESTION 268
Drag and Drop Question
You have an Azure Synapse Analytics SQL pool named Pool1 on a logical Microsoft SQL server named Server1.
You need to implement Transparent Data Encryption (TDE) on Pool1 by using a custom key named key1.
Which five actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to

the answer area and arrange them in the correct order.

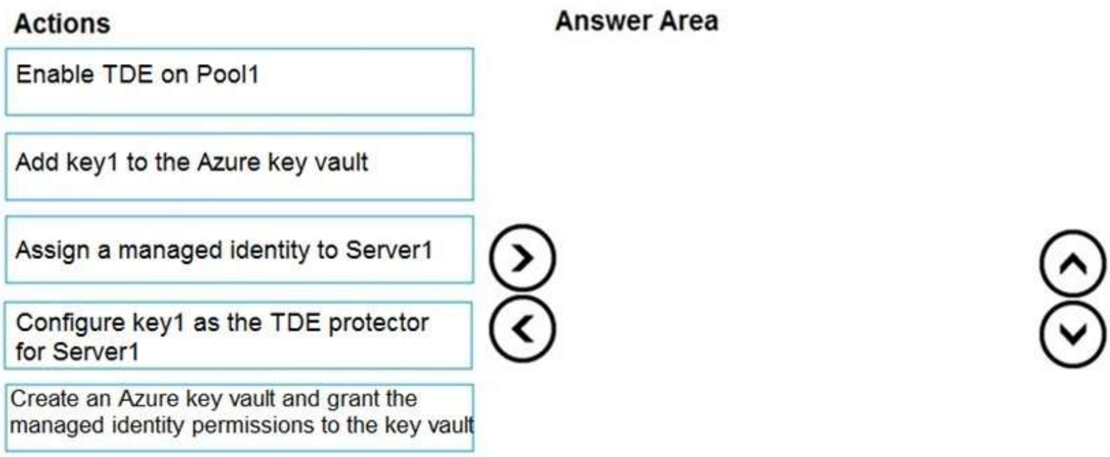

Actions

- Enable TDE on Pool1
- Add key1 to the Azure key vault
- Assign a managed identity to Server1
- Configure key1 as the TDE protector for Server1
- Create an Azure key vault and grant the managed identity permissions to the key vault

Answer Area

Answer:

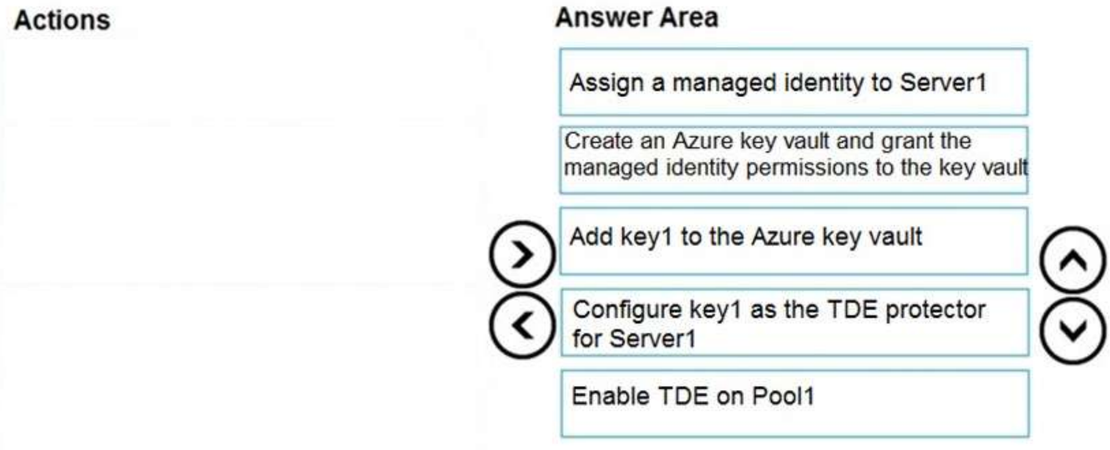

Actions

Answer Area

1. Assign a managed identity to Server1
2. Create an Azure key vault and grant the managed identity permissions to the key vault
3. Add key1 to the Azure key vault
4. Configure key1 as the TDE protector for Server1
5. Enable TDE on Pool1

Explanation:
Step 1: Assign a managed identity to Server1
You will need an existing Managed Instance as a prerequisite.
Step 2: Create an Azure key vault and grant the managed identity permissions to the vault Create Resource and setup Azure Key Vault.
Step 3 :Add key1 to the Azure key vault
The recommended way is to import an existing key from a .pfx file or get an existing key from the vault. Alternatively, generate a new key directly in Azure Key Vault.
Step 4: Configure key1 as the TDE protector for Server1
Provide TDE Protector key
Step 5: Enable TDE on Pool1
Reference:
https://docs.microsoft.com/en-us/azure/azure-sql/managed-instance/scripts/transparent-data-encryption-byok-powershell

QUESTION 269
Drag and Drop Question
You have an Azure Active Directory (Azure AD) tenant that contains a security group named Group1. You have an Azure Synapse Analytics dedicated SQL pool named dw1 that contains a schema named schema1.
You need to grant Group1 read-only permissions to all the tables and views in schema1. The solution must use the principle of least privilege.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

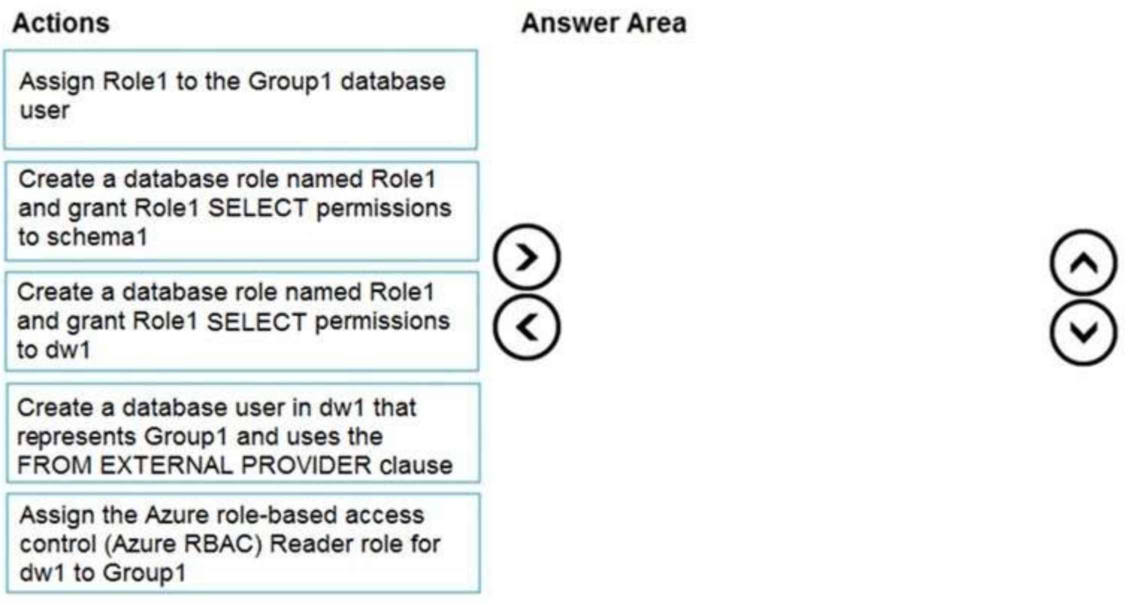

Actions

- Assign Role1 to the Group1 database user
- Create a database role named Role1 and grant Role1 SELECT permissions to schema1
- Create a database role named Role1 and grant Role1 SELECT permissions to dw1
- Create a database user in dw1 that represents Group1 and uses the FROM EXTERNAL PROVIDER clause
- Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1

Answer Area

Answer:

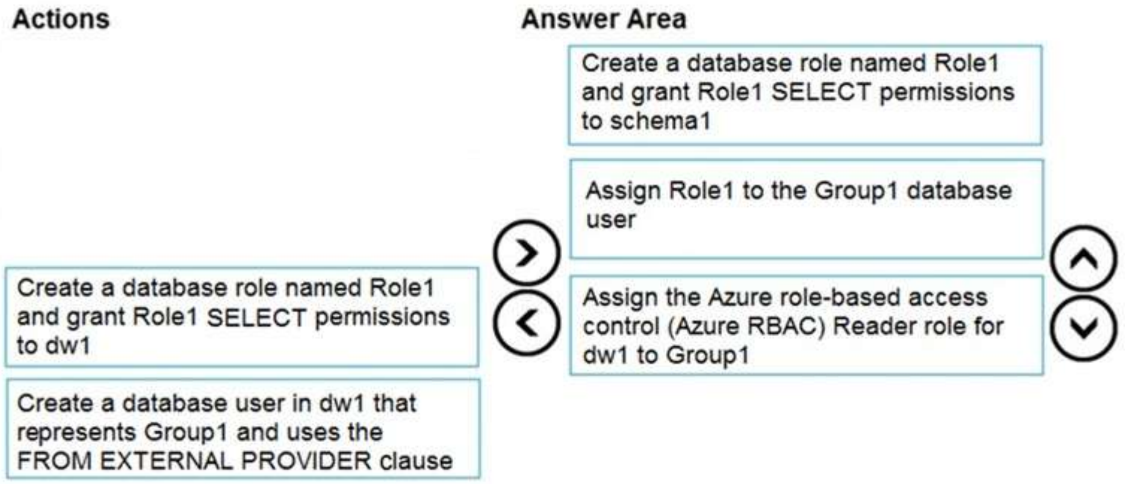

Actions

- Create a database role named Role1 and grant Role1 SELECT permissions to dw1
- Create a database user in dw1 that represents Group1 and uses the FROM EXTERNAL PROVIDER clause

Answer Area

- Create a database role named Role1 and grant Role1 SELECT permissions to schema1
- Assign Role1 to the Group1 database user
- Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1

Explanation:
Step 1: Create a database role named Role1 and grant Role1 SELECT permissions to schema You need to grant Group1 read-only permissions to all the tables and views in schema1. Place one or more database users into a database role and then assign permissions to the database role.
Step 2: Assign Rol1 to the Group database user
Step 3: Assign the Azure role-based access control (Azure RBAC) Reader role for dw1 to Group1
Reference:
https://docs.microsoft.com/en-us/azure/data-share/how-to-share-from-sql

QUESTION 270
Hotspot Question
You have an Azure subscription that contains the following resources:

```
- An Azure Active Directory (Azure AD) tenant that contains a security group named
Group1
- An Azure Synapse Analytics SQL pool named Pool1
```

You need to control the access of Group1 to specific columns and rows in a table in Pool1.
Which Transact-SQL commands should you use? To answer, select the appropriate options in the answer area.

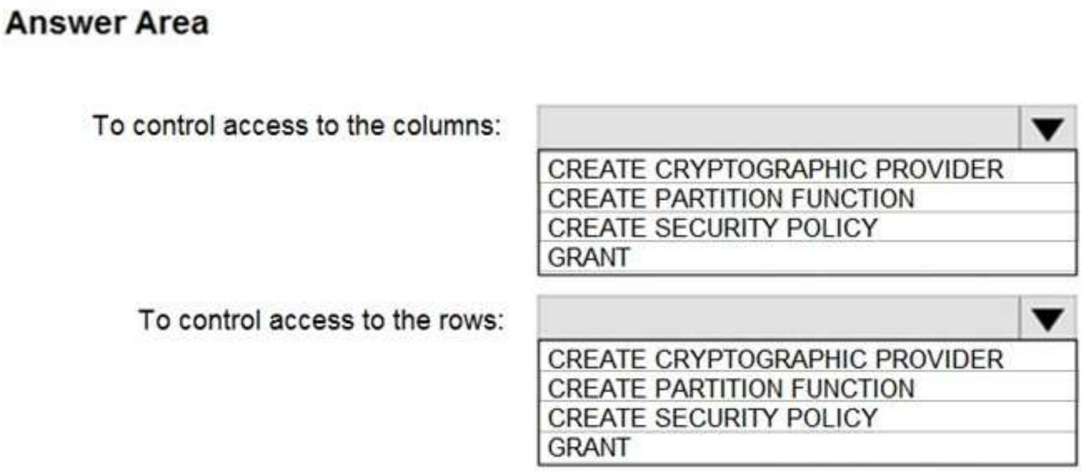

Answer Area

To control access to the columns:

- CREATE CRYPTOGRAPHIC PROVIDER
- CREATE PARTITION FUNCTION
- CREATE SECURITY POLICY
- GRANT

To control access to the rows:

- CREATE CRYPTOGRAPHIC PROVIDER
- CREATE PARTITION FUNCTION
- CREATE SECURITY POLICY
- GRANT

Answer:

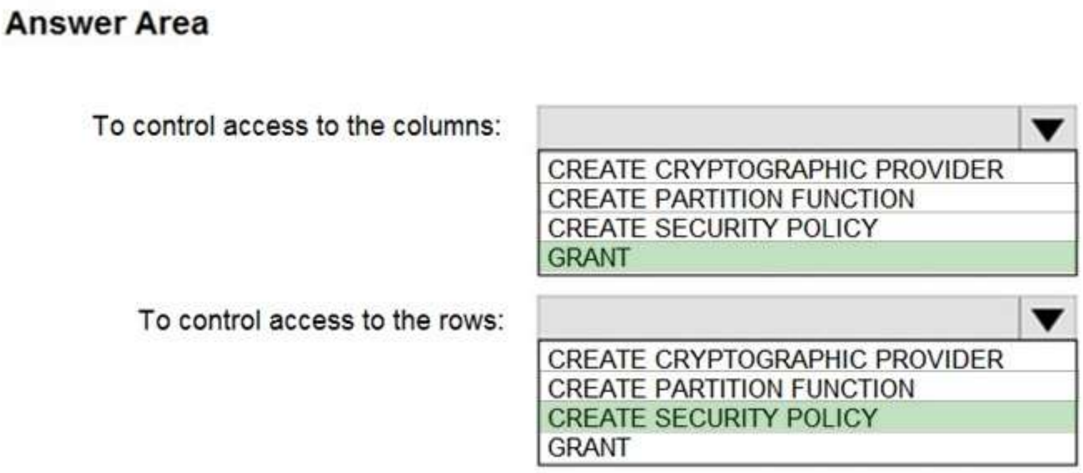

Answer Area

To control access to the columns:

- CREATE CRYPTOGRAPHIC PROVIDER
- CREATE PARTITION FUNCTION
- CREATE SECURITY POLICY
- [✓ GRANT]

To control access to the rows:

- CREATE CRYPTOGRAPHIC PROVIDER
- CREATE PARTITION FUNCTION
- [✓ CREATE SECURITY POLICY]
- GRANT

Explanation:
Box 1: GRANT
You can implement column-level security with the GRANT T-SQL statement.
Box 2: CREATE SECURITY POLICY
Implement Row Level Security by using the CREATE SECURITY POLICY Transact-SQL statement
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/column-level-security

QUESTION 271
Hotspot Question
You have an Azure Cosmos DB database.
You need to use Azure Stream Analytics to check for uneven distributions of queries that can affect performance.
Which two settings should you configure? To answer, select the appropriate settings in the answer area.
NOTE: Each correct selection is worth one point.

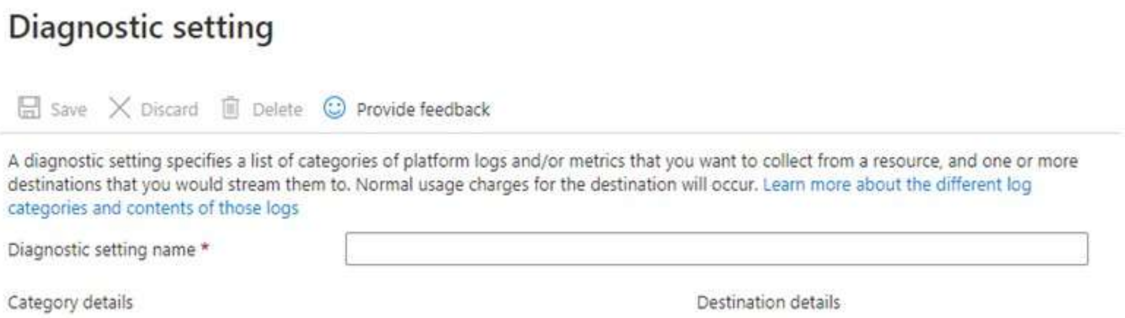

Diagnostic setting

A diagnostic setting specifies a list of categories of platform logs and/or metrics that you want to collect from a resource, and one or more destinations that you would stream them to. Normal usage charges for the destination will occur. Learn more about the different log categories and contents of those logs

Diagnostic setting name *

Category details

Destination details

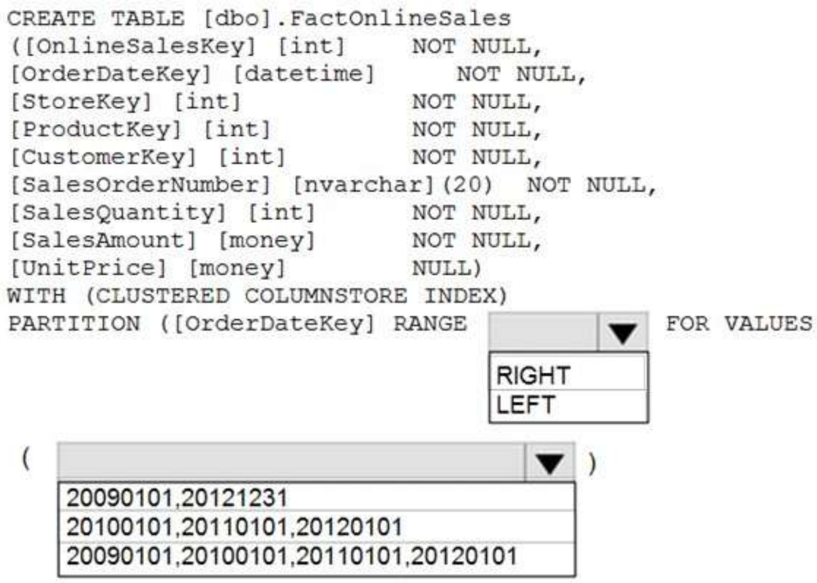

```
CREATE TABLE [dbo].FactOnlineSales
([OnlineSalesKey] [int]     NOT NULL,
[OrderDateKey] [datetime]      NOT NULL,
[StoreKey] [int]            NOT NULL,
[ProductKey] [int]          NOT NULL,
[CustomerKey] [int]         NOT NULL,
[SalesOrderNumber] [nvarchar](20)   NOT NULL,
[SalesQuantity] [int]       NOT NULL,
[SalesAmount] [money]       NOT NULL,
[UnitPrice] [money]         NULL)
WITH (CLUSTERED COLUMNSTORE INDEX)
PARTITION ([OrderDateKey] RANGE [          ▼] FOR VALUES
```

```
 ( [                              ▼] )
```

Answer:

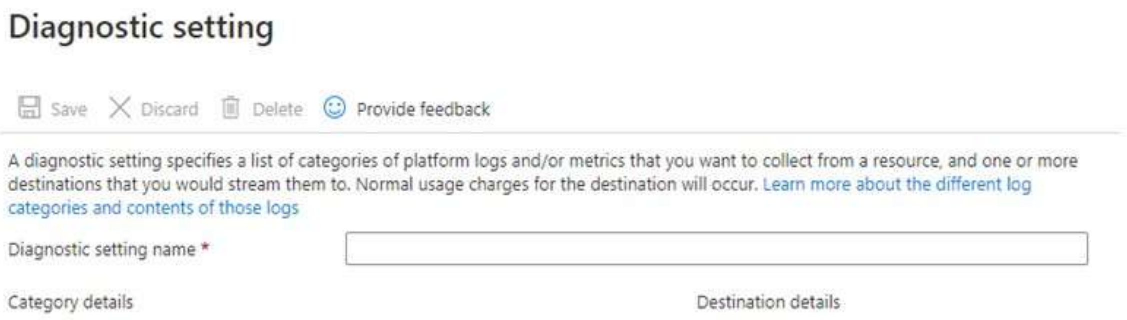

Diagnostic setting

A diagnostic setting specifies a list of categories of platform logs and/or metrics that you want to collect from a resource, and one or more destinations that you would stream them to. Normal usage charges for the destination will occur. Learn more about the different log categories and contents of those logs

Diagnostic setting name *

Category details

Destination details

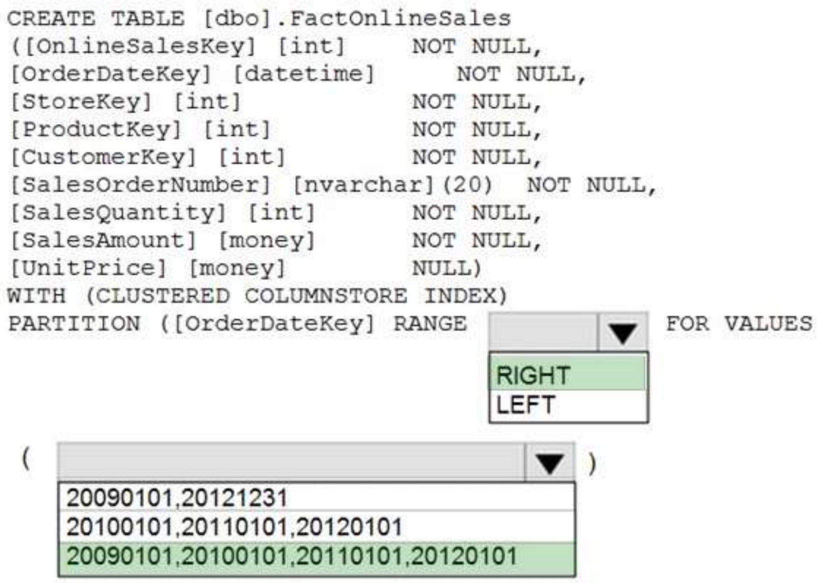

```
CREATE TABLE [dbo].FactOnlineSales
([OnlineSalesKey] [int]     NOT NULL,
[OrderDateKey] [datetime]      NOT NULL,
[StoreKey] [int]            NOT NULL,
[ProductKey] [int]          NOT NULL,
[CustomerKey] [int]         NOT NULL,
[SalesOrderNumber] [nvarchar](20)  NOT NULL,
[SalesQuantity] [int]       NOT NULL,
[SalesAmount] [money]       NOT NULL,
[UnitPrice] [money]         NULL)
WITH (CLUSTERED COLUMNSTORE INDEX)
PARTITION ([OrderDateKey] RANGE          FOR VALUES
```

RIGHT
LEFT

( )

20090101,20121231
20100101,20110101,20120101
20090101,20100101,20110101,20120101

Explanation:
Box 1: RIGHT
Use right for dates.
1- RIGHT means < or >=
2- LEFT means <= and >.
Box 2: 20090101, 201001010, 20110101, 20120101
Four values are better than three or two.

QUESTION 272
Hotspot Question
You have an Azure Cosmos DB database.
You need to use Azure Stream Analytics to check for uneven distributions of queries that can affect performance.
Which two settings should you configure? To answer, select the appropriate settings in the answer area.
NOTE: Each correct selection is worth one point.

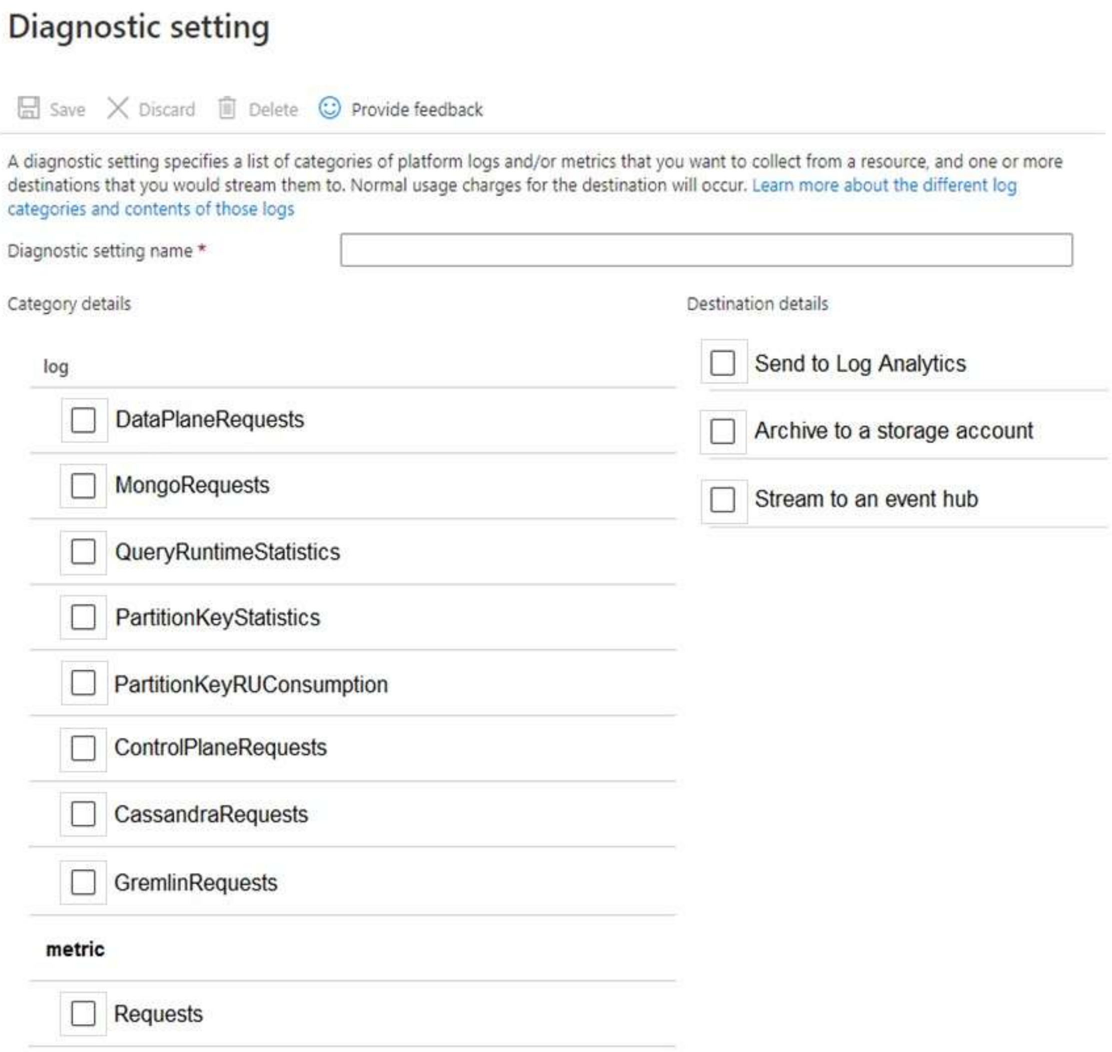

Diagnostic setting

A diagnostic setting specifies a list of categories of platform logs and/or metrics that you want to collect from a resource, and one or more destinations that you would stream them to. Normal usage charges for the destination will occur. Learn more about the different log categories and contents of those logs

Diagnostic setting name *

Category details

log

□ DataPlaneRequests

□ MongoRequests

□ QueryRuntimeStatistics

□ PartitionKeyStatistics

□ PartitionKeyRUConsumption

□ ControlPlaneRequests

□ CassandraRequests

□ GremlinRequests

metric

□ Requests

Destination details

□ Send to Log Analytics

□ Archive to a storage account

□ Stream to an event hub

Answer:

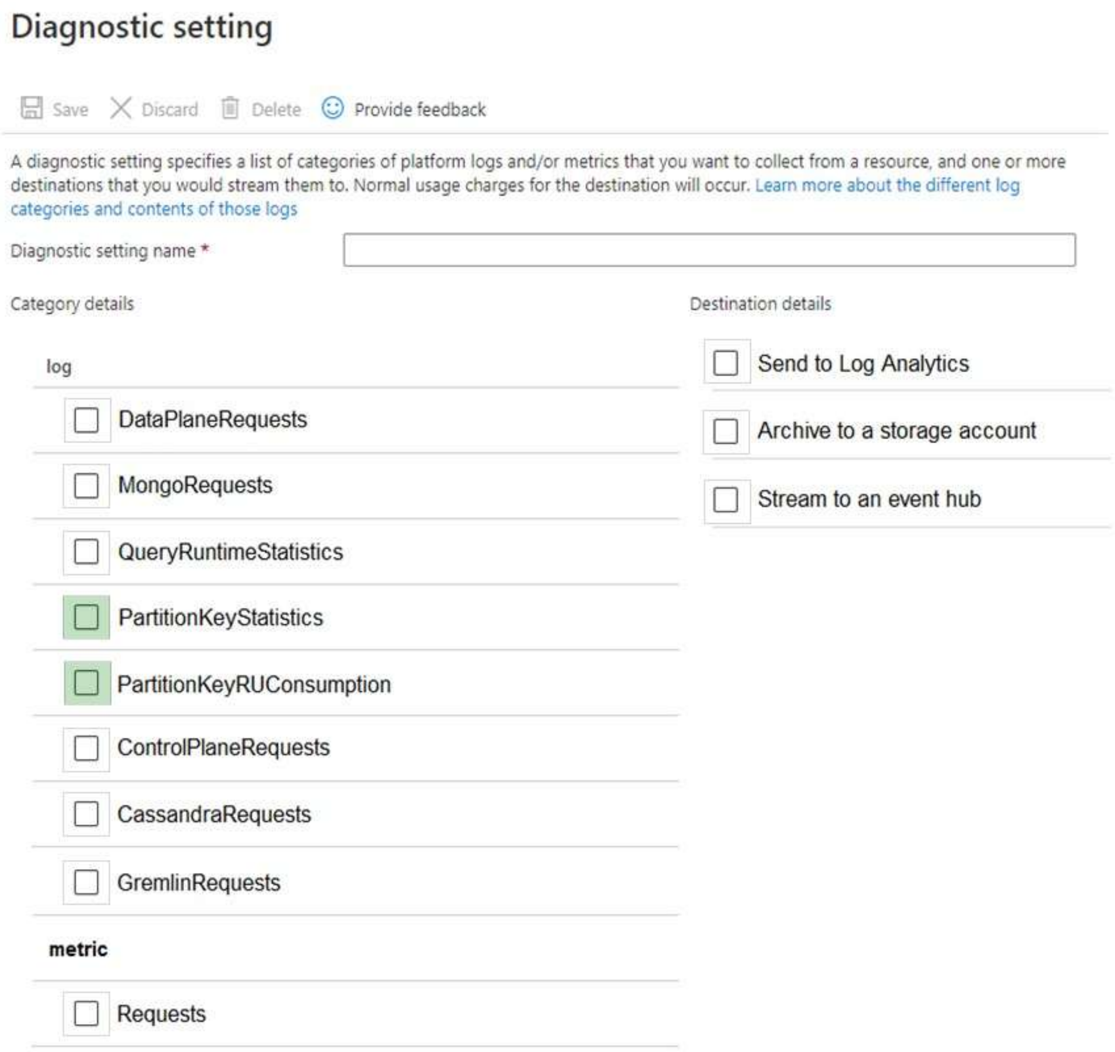

Diagnostic setting

A diagnostic setting specifies a list of categories of platform logs and/or metrics that you want to collect from a resource, and one or more destinations that you would stream them to. Normal usage charges for the destination will occur. Learn more about the different log categories and contents of those logs

Diagnostic setting name *

Category details

log

□ DataPlaneRequests

□ MongoRequests

□ QueryRuntimeStatistics

☑ PartitionKeyStatistics

☑ PartitionKeyRUConsumption

□ ControlPlaneRequests

□ CassandraRequests

□ GremlinRequests

metric

□ Requests

Destination details

□ Send to Log Analytics

□ Archive to a storage account

□ Stream to an event hub

Explanation:
PartitionKeyStatistics: Select this option to log the statistics of the partition keys. This is currently represented with the storage size (KB) of the partition keys.
PartitionKeyRUConsumption: This log reports the aggregated per-second RU/s consumption of partition keys. Currently, Azure Cosmos DB reports partition keys for SQL API accounts only and for point read/ write and stored procedure operations. other APIs and operation types are not supported. For other APIs, the partition key column in the diagnostic log table will be empty. This log contains data such as subscription ID, region name, database name, collection name, partition key, operation type, and request charge.
Note:
How to get partition key statistics to evaluate skew across top 3 partitions for a database account:
AzureDiagnostics
| where ResourceProvider=="MICROSOFT.DOCUMENTDB" and Category=="PartitionKeyStatistics" | project SubscriptionId, regionName_s, databaseName_s, collectionName_s, partitionKey_s, sizeKb_d, ResourceId
Incorrect Answers:
DataPlaneRequests: Select this option to log back-end requests to the SQL API accounts in Azure Cosmos DB.
Reference:
https://docs.microsoft.com/en-us/azure/cosmos-db/cosmosdb-monitor-resource-logs