

Databricks

*Azure-Databricks-Certified-Associate-Platform-Administrator
Azure Databricks Certified Associate Platform Administrator Exam*

- **Up to Date products, reliable and verified.**
- **Questions and Answers in PDF Format.**

Full Version Features:

- **90 Days Free Updates**
- **30 Days Money Back Guarantee**
- **Instant Download Once Purchased**
- **24 Hours Live Chat Support**

For More Information:

<https://www.testsexpert.com/>

• **Product Version**

Latest Version: 10.0

Question: 1

A user have the below permission, can he be able to clone an existing cluster ? (Choose the least privileged)

- A. No permissions permission
- B. Can attach to
- C. Can restart
- D. Can manage

Answer: A

Explanation:

a user with no permissions on the cluster but have cluster creation privilege can able to clone the cluster (without can manage ,can restart, can attach)

Question: 2


IP Access lists have configured with azure databricks workspace to allow/block ip address range . A User wants to access the databricks workspace , which of the following scenario's he can able to access ?

- A. User's IP Address configured in Block list but not the Allow list
- B. User's IP Address is not configured in Block list , but there are zero allow lists for workspace.
- C. User's IP address configured in Allow list but not in Block list
- D. User's IP address configured in both Block list and Allow list

Answer: B,C

Explanation:

IP access list flow



For additional documentation :

<https://docs.microsoft.com/en-us/azure/databricks/security/network/ip-access-list>

Question: 3

Which of the following statements regarding permissions are correct ?

- A. Any admin can assign other users as admins in the Admin console
- B. There is only one way to grant users the ability to create clusters in the admin console
- C. Any one can access admin console in the Azure Databricks UI.

Answer: A

Question: 4

When you mount data using a cluster enabled with Azure Data Lake Storage credential passthrough, any read or write to the mount point uses your Azure AD credentials. This mount point will be visible to other users, but the only users that will have read and write access are those who:

- A. Are using a cluster disabled for Azure Data Lake Storage credential passthrough
- B. Have access to the underlying Azure Data Lake Storage storage account
- C. Are using a cluster enabled for Azure Data Lake Storage credential passthrough
- D. Don't have access to the underlying Azure Data Lake Storage storage account

Answer: B,C

Explanation:

When you mount data using a cluster enabled with Azure Data Lake Storage credential passthrough, any read or write to the mount point uses your Azure AD credentials. This mount point will be visible to other users, but the only users that will have read and write access are those who:

Have access to the underlying Azure Data Lake Storage storage account

Are using a cluster enabled for Azure Data Lake Storage credential passthrough

For additional documentation on credential passthrough : <https://docs.microsoft.com/en-us/azure/databricks/security/credential-passthrough/adls-passthrough>

Question: 5

Which process prevents queries from monopolizing cluster resources, examines most common causes of large queries and terminates that pass a threshold ?

- A. Task kill
- B. Query watchdog
- C. Task preemption

Answer: B

Explanation:

Guardrails for High Concurrency Clusters

Query Watchdog

- Process prevents queries from monopolizing cluster resources
- Examines most common causes of large queries and terminates queries that pass a threshold
- Enabled for all all-purpose clusters created using the UI

Task Preemption

- Scheduler automatically preempts tasks to enforce fair sharing of resources
- Guarantees interactive response times on clusters with many concurrently running jobs

Question: 6

Buying a Pre purchase plan (P3) at what level makes maximum profits?

- A. Resource Group
- B. Subscription
- C. Workspace
- D. Enterprise agreement

Answer: D

Question: 7

How to ensure consistent application of workspace tags ?

- A. Azure Policies
- B. Azure databricks policies REST API
- C. Databricks cluster policies

Answer: A

Explanation:

Leverage Azure Policies to ensure consistent application of workspace tags.

Leverage Databricks cluster policies to ensure consistent application of cluster tags.

Chargebacks

- Create separate workspaces for different divisions/teams
 - Associate workspace tag with each workspace
 - Leverage Azure Policies to ensure consistent application of workspace tags
- Create separate clusters for each project/initiative within a division/team
 - Associate cluster tags with each project
 - Leverage Databricks Cluster Policies to ensure consistent application of cluster tags
- Monitor & report chargebacks by workspace/cluster tags



Question: 8

What does Databricks Light has ?

- A. Native vectorized query engine on Databricks
- B. used to create a cluster optimized for machine learning
- C. Databricks packaging of the open source Apache Spark runtime

Answer: C

Explanation:

Photon Runtime is the Native vectorized query engine on Databricks , and Databricks Runtime for Machine Learning (Databricks Runtime ML) automates the creation of a cluster optimized for machine learning

Databricks Light has Databricks packaging of the open source Apache Spark runtime

Question: 9

Which of the below is true with using Autoscaling enabled clusters for Streaming workloads ?

- A. The default behavior with autoscaling on Structured Streaming is that the cluster will always scale up to the max number of nodes.
- B. The default behavior with autoscaling on Structured Streaming is that the cluster will always scale up to the min number of nodes.

Answer: A

Explanation:

Do not use autoscaling for streaming workloads. The default behavior with autoscaling on Structured Streaming is that the cluster will always scale up to the max number of nodes.

The recommended best practice is to disable autoscaling for Structured Streaming workloads and run it as a Databricks job on a new jobs cluster with infinite retries

Question: 10

How to ensure consistent application of cluster tags ?

- A. Databricks cluster policies
- B. azure databricks policies API
- C. Azure policies

Answer: A

Explanation:

Leverage Azure Policies to ensure consistent application of workspace tags.

Leverage Databricks cluster policies to ensure consistent application of cluster tags.

Chargebacks

- Create separate workspaces for different divisions/teams
 - Associate workspace tag with each workspace
 - Leverage Azure Policies to ensure consistent application of workspace tags
- Create separate clusters for each project/initiative within a division/team
 - Associate cluster tags with each project
 - Leverage Databricks Cluster Policies to ensure consistent application of cluster tags
- Monitor & report chargebacks by workspace/cluster tags



Question: 11

Automatic termination related to which of the below ?

- A. Inactive cluster
- B. Job cluster
- C. Query watchdog

Answer: A

Explanation:

Query Watchdog is a process that prevents queries from monopolizing cluster resources by examining the most common causes of large queries and terminating queries that pass a threshold. This article describes how to enable and configure Query Watchdog.

During cluster creation, you can specify an inactivity period in minutes after which you want the cluster to terminate. If the difference between the current time and the last command run on the cluster is more than the inactivity period specified, Azure Databricks automatically terminates that cluster.

Question: 12

As an administrator , you have created a Secret scope for analytics team where they are storing azure synapse analytics credentials , Now another subgroup requested for a scope to store Azure Blob storage credentials .

Which of the below you recommend as an administrator?

- A. Create two Azure Key Vault-backed scopes with each scope referencing to one Azure Key Vault and add your secrets to same Azure Key Vault
- B. Create two Azure Key Vault-backed scopes with each scope referencing a different Azure Key Vault and add your secrets to those two Azure Key Vaults
- C. Create a Azure backed secret scope and store both of the sub groups credentials to it (by storing the secrets in same keyvault)
- D. Create a databricks secret scope and store both of the sub groups credentials to it

Answer: B

Explanation:

If you use two Azure Key Vault-backed scopes with both scopes referencing the same Azure Key Vault and add your secrets to that Azure Key Vault, all Azure Synapse Analytics and Azure Blob storage secrets will be available. Since ACLs are at the scope level, all members across the two subgroups will see all secrets.

Creating two Azure Key Vault-backed scopes with each scope referencing a different Azure Key Vault and add your secrets to those two Azure Key Vaults will restrict access to a set of secrets to each group.

Question: 13

Which of the below are true with respect to Spot VM's

- A. For non-mission critical jobs , use Driver on-demand and workers spot
- B. For workflows with tight SLA's , Use spot instance w/fallback to on-demand
- C. For Production jobs , Use on-demand instances

D. We can use spot for driver

Answer: A,B,C

Explanation:

Spot VMs (Private Preview)

- Use spot instances to use spare VM instances for below market rate
 - Great for ad-hoc/shared clusters
 - Not recommended for production
 - Never use for driver!
- Combine on-demand and spot instances (with custom spot price) to tailor clusters to different use cases

SLA	Spot or On-Demand
Non-mission critical jobs	Driver on-demand and workers spot
Workflows with tight SLAs	Use spot instance w/fallback to on-demand
Production jobs	Use on-demand instances

Should not use Spot for driver since driver nodes loses the cached data if spot instances were withdrawn.

Question: 14

With respect to Cluster permissions ,a user has only can attach permission, will he be able to see spark driver logs of the cluster?

- A. No , he cannot see the Spark driver logs since he don't have Can manage.
- B. Yes , he can able to see the spark drive logs.

Answer: B

Explanation:

Ability	No Permissions	Can Attach To	Can Restart	Can Manage
Attach notebook to cluster		x	x	x
View Spark UI		x	x	x
View cluster metrics		x	x	x
View driver logs		x	x	x
Terminate cluster			x	x
Start cluster			x	x
Restart cluster			x	x
Edit cluster				x
Attach library to cluster				x
Resize cluster				x
Modify permissions				x

Question: 15

which of the options are correct based below scenarios

- A. Use High concurrency for Analytics or BI ,Ad-hoc development
- B. Use single node for heavyload ML workloads
- C. Use Standard cluster for production batch or streaming ETL/ML jobs , Ad-hoc development

Answer: A,C

Explanation:

Cluster Modes - Use Cases

	Standard	High Concurrency	Single Node
Use Cases	Production batch or streaming ETL/ML jobs, Ad-hoc development	Analytics or BI, Ad-hoc development	Lightweight EDA, Single node ML workloads
Audience	Data engineer or data scientist	Data analysts, management, finance	Data Analyst or Data Scientist

Instance Types

Workload Type	Azure Type	Recommended Use Case
Memory Optimized	DSv2	Memory-intensive applications
Compute Optimized	Fsv2	Structured Streaming, Distributed Analytics, Data Science Applications
Storage Optimized	Lsv2	Use cases that require higher disk throughput and IO
General Purpose	DSv2	Enterprise-grade applications, relational databases, and analytics with in-memory caching



For More Information – **Visit link below:**
<https://www.testsexpert.com/>

Features:

■ Money Back Guarantee.....



■ 100% Course Coverage.....



■ 90 Days Free Updates.....



■ Instant Email Delivery after Order.....

