



# 70-775<sup>Q&As</sup>

Perform Data Engineering on Microsoft Azure HDInsight

## Pass Microsoft 70-775 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/70-775.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers






### QUESTION 1

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

You are planning a big data infrastructure by using an Apache Spark cluster in Azure HDInsight. The cluster has 24 processor cores and 512 GB of memory.

The architecture of the infrastructure is shown in the exhibit. (Click the Exhibit button.)



The architecture will be used by the following users:

The data sources in the batch layer share a common storage container. The following data sources are used:

The business analysts report that they experience performance issues when they run the monitoring queries.

You troubleshoot the performance issues and discover that the intermediate tables generated when the analysts run the queries cause pressure for the Java Virtual Machine (JVM) garbage collection per job.

Which configuration settings should you modify to alleviate the performance issues?

- A. `spark.sql.inMemoryColumnarStorage.batchSize`
- B. `spark.sql.broadcastTimeout`
- C. `spark.sql.files.openCostInBytes`
- D. `spark.sql.shuffle.partitions`

Correct Answer: D

### QUESTION 2



You have a domain-joined Azure HDInsight cluster.

You plan to assign permissions to several support staff.

You need to assign roles to the staff so that they can perform specific tasks. The solution must use the principle of least privilege.

Which role should you assign for each task? To answer, drag the appropriate roles to the correct targets. Each role may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Select and Place:

### Roles

- Cluster Administrator
- Cluster Operator
- Cluster User
- Service Administrator
- Service Operator

### Answer Area

#### Task

- Configure services
- View service-level alerts:
- View cluster configurations:

### Roles

- Role
- Role
- Role

Correct Answer:

### Roles

- Cluster Administrator
- Cluster Operator
- Cluster User
- Service Administrator
- Service Operator

### Answer Area

#### Task

- Configure services
- View service-level alerts:
- View cluster configurations:

### Roles

- Service Administrator
- Cluster User
- Cluster User

### QUESTION 3

You are evaluating the use of Azure HDInsight clusters for various workloads.

Which type of HDInsight cluster should you create for each workload? To answer, drag the appropriate cluster types to the correct workloads. Each cluster type may be used once, more than once, or not at all. You may need to drag the split



bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

### Cluster types

- Apache Hadoop
- Apache Spark
- Apache HBase
- Apache Storm

### Answer Area

- ETL/ELT:
- Data in motion.
- Transactional processing:

Cluster type

Cluster type

Cluster type

Correct Answer:

### Cluster types

- Apache Spark

### Answer Area

- ETL/ELT:
- Data in motion.
- Transactional processing:

Apache Hadoop

Apache Storm

Apache HBase

## QUESTION 4

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are building a security tracking solution in Apache Kafka to parse security logs. The security logs record an entry each time a user attempts to access an application. Each log entry contains the IP address used to make the attempt and

the country from which the attempt originated.



You need to receive notifications when an IP address from outside of the United States is used to access the application.

Solution: Create two new brokers. Create a file import process to send messages. Run the producer.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

---

### QUESTION 5

You have an Azure HDInsight cluster.

You need to store data in a file format that maximizes compression and increases read performance.

Which type of file format should you use?

A. ORC

B. Apache Parquet

C. Apache Avro

D. Apache Sequence

Correct Answer: A

References: <http://www.semantikoz.com/blog/orc-intelligent-big-data-file-format-hadoop-hive/>

[70-775 Study Guide](#)

[70-775 Exam Questions](#)

[70-775 Braindumps](#)





To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

## Try our product !


- 100% Guaranteed Success
- 100% Money Back Guarantee
- 365 Days Free Update
- Instant Download After Purchase
- 24x7 Customer Support
- Average 99.9% Success Rate
- More than 800,000 Satisfied Customers Worldwide
- Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.geekcert.com/allproducts>

## Need Help

Please provide as much detail as possible so we can best assist you.  
To update a previously submitted ticket:



 <p><b>One Year Free Update</b> Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p><b>Money Back Guarantee</b> To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p><b>Security &amp; Privacy</b> We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information &amp; peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.  
All trademarks are the property of their respective owners.  
Copyright © geekcert, All Rights Reserved.