



# 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.pass4lead.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**

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 2**

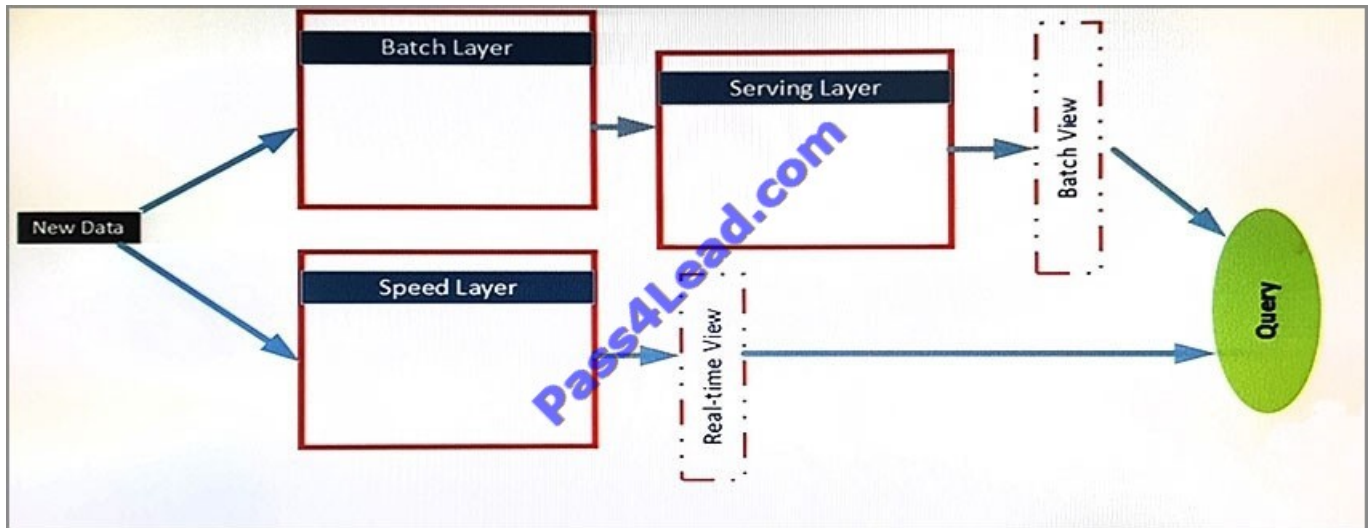
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is



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:

You need to ensure that the analysts can query the logistics data by using JDBC APIs and SQL APIs.

Which technology should you implement?

- A. Apache Phoenix
- B. Apache Spark
- C. Apache Storm
- D. Apache Hive

Correct Answer: A

References: <http://phoenix.apache.org/>

### QUESTION 3

You plan to copy data from Azure Blob storage to an Azure SQL database by using Azure Data Factory. Which file formats can you use?

- A. binary, JSON, Apache Parquet, and ORC
- B. OXPS, binary, text and JSON



- C. XML, Apache Avro, text, and ORC
- D. text, JSON, Apache Avro, and Apache Parquet

Correct Answer: D


References: <https://docs.microsoft.com/en-us/azure/data-factory/supported-file-formats-and-compression-codecs>

#### QUESTION 4

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 5

You use YARN to manage the resources for a Spark Thrift Server running on a Linux- based Apache Spark cluster in Azure HDInsight.

You discover that the cluster does not fully utilize the resources. You want to increase resource allocation.

You need to increase the number of executors and the allocation of memory to the Spark Thrift Server driver.

Which two parameters should you modify? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. spark.dynamicAllocation.maxExecutors

B. spark.cores.max

C. spark.executor.memory

D. spark\_thrift\_cmd\_opts

E. spark.executor.instances

Correct Answer: AC

References: <https://stackoverflow.com/questions/37871194/how-to-tune-spark-executor-number-cores-and-executor-memory>

[70-775 Practice Test](#)

[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.pass4lead.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 © pass4lead, All Rights Reserved.