



Vendor: Microsoft

Exam Code: AZ-305

Exam Name: Designing Microsoft Azure Infrastructure
Solutions

Version: DEMO

QUESTION 1

Case Study 1 - Litware

Existing Environment

Azure Environment

Litware has 10 Azure subscriptions that are linked to the Litware.com tenant and five Azure subscriptions that are linked to the dev.litware.com tenant. All the subscriptions are in an Enterprise Agreement (EA).

The litware.com tenant contains a custom Azure role-based access control (Azure RBAC) role named Role1 that grants the DataActions read permission to the blobs and files in Azure Storage.

On-Premises Environment

The on-premises network of Litware contains the resources shown in the following table.

Name	Type	Configuration
SERVER1 SERVER2 SERVER3	Ubuntu 18.04 virtual machines hosted on Hyper-V	The virtual machines host a third-party app named App1. App1 uses an external storage solution that provides Apache Hadoop-compatible data storage. The data storage supports POSIX access control list (ACL) file-level permissions.
SERVER10	Server that runs Windows Server 2016	The server contains a Microsoft SQL Server instance that hosts two databases named DB1 and DB2.

Network Environment

Litware has ExpressRoute connectivity to Azure.

Planned Changes and Requirements

Litware plans to implement the following changes:

- Migrate DB1 and DB2 to Azure.
- Migrate App1 to Azure virtual machines.
- Migrate the external storage used by App1 to Azure Storage.
- Deploy the Azure virtual machines that will host App1 to Azure dedicated hosts.

Hotspot Question

You plan to migrate App1 to Azure.

You need to recommend a storage solution for App1 that meets the security and compliance requirements.

Which type of storage should you recommend, and how should you recommend configuring the storage? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Storage account type:

	▼
Premium page blobs	
Premium file shares	
Standard general-purpose v2	

Configuration:

	▼
NFSv3	
Large file shares	
Hierarchical namespace	

Answer:

Answer Area

Storage account type:

	▼
Premium page blobs	
Premium file shares	
Standard general-purpose v2	

Configuration:

	▼
NFSv3	
Large file shares	
Hierarchical namespace	

Explanation:

Box 1: Standard general-purpose v2

Standard general-purpose v2 supports Blob Storage.

Azure Storage provides data protection for Blob Storage and Azure Data Lake Storage Gen2.

Box 2: Hierarchical namespace

Scenario: Plan: Migrate App1 to Azure virtual machines.

Azure Data Lake Storage Gen2 implements an access control model that supports both Azure role-based access control (Azure RBAC) and POSIX-like access control lists (ACLs). Data Lake Storage Gen2 and the Network File System (NFS) 3.0 protocol both require a storage account with a hierarchical namespace enabled.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-protection-overview>

<https://docs.microsoft.com/en-us/azure/storage/blobs/immutable-storage-overview>

QUESTION 2

Case Study 2 - Fabrikam, Inc

Overview

Fabrikam, Inc. is an engineering company that has offices throughout Europe. The company has a main office in London and three branch offices in Amsterdam, Berlin, and Rome.

Existing Environment: Active Directory Environment

The network contains two Active Directory forests named corp.fabrikam.com and rd.fabrikam.com. There are no trust relationships between the forests.

Corp.fabrikam.com is a production forest that contains identities used for internal user and computer authentication.

Rd.fabrikam.com is used by the research and development (R&D) department only. The R&D department is restricted to using on-premises resources only.

Existing Environment: Network Infrastructure

Each office contains at least one domain controller from the corp.fabrikam.com domain. The main office contains all the domain controllers for the rd.fabrikam.com forest.

All the offices have a high-speed connection to the internet.

An existing application named WebApp1 is hosted in the data center of the London office. WebApp1 is used by customers to place and track orders. WebApp1 has a web tier that uses Microsoft Internet Information Services (IIS) and a database tier that runs Microsoft SQL Server 2016. The web tier and the database tier are deployed to virtual machines that run on Hyper-V.

The IT department currently uses a separate Hyper-V environment to test updates to WebApp1.

Fabrikam purchases all Microsoft licenses through a Microsoft Enterprise Agreement that includes Software Assurance.

Existing Environment: Problem Statements

The use of WebApp1 is unpredictable. At peak times, users often report delays. At other times, many resources for WebApp1 are underutilized.

You need to recommend a strategy for migrating the database content of WebApp1 to Azure. What should you include in the recommendation?

- A. Use Azure Site Recovery to replicate the SQL servers to Azure.
- B. Use SQL Server transactional replication.
- C. Copy the BACPAC file that contains the Azure SQL database file to Azure Blob storage.
- D. Copy the VHD that contains the Azure SQL database files to Azure Blob storage

Answer: D

Explanation:

Before you upload a Windows virtual machine (VM) from on-premises to Azure, you must prepare the virtual hard disk (VHD or VHDX).

Scenario: WebApp1 has a web tier that uses Microsoft Internet Information Services (IIS) and a database tier that runs Microsoft SQL Server 2016. The web tier and the database tier are deployed to virtual machines that run on Hyper-V.

Reference:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/prepare-for-upload-vhd-image>

QUESTION 3

Case Study 3 - Contoso

Existing Environment: Technical Environment

The on-premises network contains a single Active Directory domain named contoso.com. Contoso has a single Azure subscription.

Existing Environment: Business Partnerships

Contoso has a business partnership with Fabrikam, Inc. Fabrikam users access some Contoso applications over the internet by using Azure Active Directory (Azure AD) guest accounts.

Requirements: Planned Changes

Contoso plans to deploy two applications named App1 and App2 to Azure.

Requirements: App1

App1 will be a Python web app hosted in Azure App Service that requires a Linux runtime. Users from Contoso and Fabrikam will access App1.

You need to recommend a solution that meets the data requirements for App1.

What should you recommend deploying to each availability zone that contains an instance of App1?

- A. an Azure Cosmos DB that uses multi-region writes
- B. an Azure Data Lake store that uses geo-zone-redundant storage (GZRS)
- C. an Azure SQL database that uses active geo-replication
- D. an Azure Storage account that uses geo-zone-redundant storage (GZRS)

Answer: A

Explanation:

Scenario: App1 has the following data requirements:

Each instance will write data to a data store in the same availability zone as the instance.

Data written by any App1 instance must be visible to all App1 instances.

Azure Cosmos DB: Each partition across all the regions is replicated. Each region contains all the data partitions of an Azure Cosmos container and can serve reads as well as serve writes when multi-region writes is enabled.

Incorrect Answers:

B, D: GZRS protects against failures. Geo-redundant storage (with GRS or GZRS) replicates your data to another physical location in the secondary region to protect against regional outages.

However, that data is available to be read only if the customer or Microsoft initiates a failover from

the primary to secondary region.

C: Active geo-replication is designed as a business continuity solution that lets you perform quick disaster recovery of individual databases in case of a regional disaster or a large scale outage. Once geo-replication is set up, you can initiate a geo-failover to a geo-secondary in a different Azure region. The geo-failover is initiated programmatically by the application or manually by the user.

Reference:

<https://docs.microsoft.com/en-us/azure/cosmos-db/high-availability>

QUESTION 4

You have an Azure Active Directory (Azure AD) tenant that syncs with an on-premises Active Directory domain.

You have an internal web app named WebApp1 that is hosted on-premises. WebApp1 uses Integrated Windows authentication.

Some users work remotely and do **NOT** have VPN access to the on-premises network.

You need to provide the remote users with single sign-on (SSO) access to WebApp1.

Which two features should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Azure AD Application Proxy
- B. Azure AD Privileged Identity Management (PIM)
- C. Conditional Access policies
- D. Azure Arc
- E. Azure AD enterprise applications
- F. Azure Application Gateway

Answer: AE

Explanation:

A: Application Proxy is a feature of Azure AD that enables users to access on-premises web applications from a remote client. Application Proxy includes both the Application Proxy service which runs in the cloud, and the Application Proxy connector which runs on an on-premises server.

You can configure single sign-on to an Application Proxy application.

E: Add an on-premises app to Azure AD

Now that you've prepared your environment and installed a connector, you're ready to add on-premises applications to Azure AD.

1. Sign in as an administrator in the Azure portal.
2. In the left navigation panel, select Azure Active Directory.
3. Select Enterprise applications, and then select New application.
4. Select Add an on-premises application button which appears about halfway down the page in the On-premises applications section. Alternatively, you can select Create your own application at the top of the page and then select Configure Application Proxy for secure remote access to an on-premise application.
5. In the Add your own on-premises application section, provide the following information about your application.
6. Etc.

Incorrect:

Not C: Conditional Access policies are not required.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/app-proxy/application-proxy-add-on-premises-application>

QUESTION 5

You have an Azure Active Directory (Azure AD) tenant named contoso.com that has a security group named Group1. Group1 is configured for assigned membership. Group1 has 50 members, including 20 guest users.

You need to recommend a solution for evaluating the membership of Group1. The solution must meet the following requirements:

- The evaluation must be repeated automatically every three months.
- Every member must be able to report whether they need to be in Group1.
- Users who report that they do not need to be in Group1 must be removed from Group1 automatically.
- Users who do not report whether they need to be in Group1 must be removed from Group1 automatically.

What should you include in the recommendation?

- A. Implement Azure AD Identity Protection.
- B. Change the Membership type of Group1 to **Dynamic User**.
- C. Create an access review.
- D. Implement Azure AD Privileged Identity Management (PIM).

Answer: C

Explanation:

Have reviews recur periodically: You can set up recurring access reviews of users at set frequencies such as weekly, monthly, quarterly or annually, and the reviewers will be notified at the start of each review. Reviewers can approve or deny access with a friendly interface and with the help of smart recommendations.

Reference:

<https://docs.microsoft.com/en-us/azure/active-directory/governance/access-reviews-overview#learn-about-access-reviews>

QUESTION 7

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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Your company deploys several virtual machines on-premises and to Azure. ExpressRoute is being deployed and configured for on-premises to Azure connectivity.

Several virtual machines exhibit network connectivity issues.

You need to analyze the network traffic to identify whether packets are being allowed or denied to the virtual machines.

Solution: Use Azure Traffic Analytics in Azure Network Watcher to analyze the network traffic.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Instead use Azure Network Watcher IP Flow Verify, which allows you to detect traffic filtering issues at a VM level.

Note: IP flow verify checks if a packet is allowed or denied to or from a virtual machine. The information consists of direction, protocol, local IP, remote IP, local port, and remote port. If the packet is denied by a security group, the name of the rule that denied the packet is returned. While any source or destination IP can be chosen, IP flow verify helps administrators quickly diagnose connectivity issues from or to the internet and from or to the on-premises environment.

Reference:

<https://docs.microsoft.com/en-us/azure/network-watcher/network-watcher-ip-flow-verify-overview>

QUESTION 8

You are developing a sales application that will contain several Azure cloud services and handle different components of a transaction. Different cloud services will process customer orders, billing, payment, inventory, and shipping.

You need to recommend a solution to enable the cloud services to asynchronously communicate transaction information by using XML messages.

What should you include in the recommendation?

- A. Azure Service Fabric
- B. Azure Data Lake
- C. Azure Service Bus
- D. Azure Traffic Manager

Answer: C

Explanation:

Asynchronous messaging options.

There are different types of messages and the entities that participate in a messaging infrastructure. Based on the requirements of each message type, Microsoft recommends Azure messaging services. The options include Azure Service Bus, Event Grid, and Event Hubs. Azure Service Bus queues are well suited for transferring commands from producers to consumers.

Data is transferred between different applications and services using messages. A message is a container decorated with metadata, and contains data. The data can be any kind of information, including structured data encoded with the common formats such as the following ones: JSON, XML, Apache Avro, Plain Text.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/guide/technology-choices/messaging>

<https://docs.microsoft.com/en-us/azure/service-bus-messaging/service-bus-messaging-overview>

QUESTION 9

You are designing a large Azure environment that will contain many subscriptions.

You plan to use Azure Policy as part of a governance solution.

To which three scopes can you assign Azure Policy definitions? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Azure Active Directory (Azure AD) administrative units
- B. Azure Active Directory (Azure AD) tenants
- C. subscriptions
- D. compute resources
- E. resource groups
- F. management groups

Answer: CEF

Explanation:

Azure Policy evaluates resources in Azure by comparing the properties of those resources to business rules. Once your business rules have been formed, the policy definition or initiative is assigned to any scope of resources that Azure supports, such as management groups, subscriptions, resource groups, or individual resources.

Reference:

<https://docs.microsoft.com/en-us/azure/governance/policy/overview>

QUESTION 10

Your company has 300 virtual machines hosted in a VMware environment. The virtual machines vary in size and have various utilization levels.

You plan to move all the virtual machines to Azure.

You need to recommend how many and what size Azure virtual machines will be required to move the current workloads to Azure. The solution must minimize administrative effort.

What should you use to make the recommendation?

- A. Azure Pricing calculator
- B. Azure Advisor
- C. Azure Migrate
- D. Azure Cost Management

Answer: C

Explanation:

<https://docs.microsoft.com/en-us/azure/migrate/migrate-appliance#collected-data---vmware>
"Metadata discovered by the Azure Migrate appliance helps you to figure out whether servers are ready for migration to Azure, right-size servers, plans costs, and analyze application dependencies".

<https://docs.microsoft.com/en-us/learn/modules/design-your-migration-to-azure/2-plan-your-azure-migration>

QUESTION 11

You plan provision a High Performance Computing (HPC) cluster in Azure that will use a third-party scheduler.

You need to recommend a solution to provision and manage the HPC cluster node.

What should you include in the recommendation?

- A. Azure Automation
- B. Azure CycleCloud
- C. Azure Purview
- D. Azure Lighthouse

Answer: B

Explanation:

You can dynamically provision Azure HPC clusters with Azure CycleCloud. Azure CycleCloud is the simplest way to manage HPC workloads.

Note: Azure CycleCloud is an enterprise-friendly tool for orchestrating and managing High Performance Computing (HPC) environments on Azure. With CycleCloud, users can provision infrastructure for HPC systems, deploy familiar HPC schedulers, and automatically scale the infrastructure to run jobs efficiently at any scale. Through CycleCloud, users can create different types of file systems and mount them to the compute cluster nodes to support HPC workloads.

Reference:

<https://docs.microsoft.com/en-us/azure/cyclecloud/overview>

QUESTION 12

You have data files in Azure Blob Storage.

You plan to transform the files and move them to Azure Data Lake Storage.

You need to transform the data by using mapping data flow.

Which service should you use?

- A. Azure Data Box Gateway
- B. Azure Databricks
- C. Azure Data Factory
- D. Azure Storage Sync

Answer: C

Explanation:

You can use Copy Activity in Azure Data Factory to copy data from and to Azure Data Lake Storage Gen2, and use Data Flow to transform data in Azure Data Lake Storage Gen2.

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-data-lake-storage>

QUESTION 13

You plan to deploy an app that will use an Azure Storage account.

You need to deploy the storage account.

The solution must meet the following requirements:

- Store the data of multiple users.
- Encrypt each user's data by using a separate key.
- Encrypt all the data in the storage account by using Microsoft keys or customer-managed keys.

What should you deploy?

- A. files in a general purpose v2 storage account.
- B. blobs in an Azure Data Lake Storage Gen2 account.
- C. files in a premium file share storage account.
- D. blobs in a general purpose v2 storage account

Answer: D

Explanation:

You can specify a customer-provided key on Blob storage operations. A client making a read or write request against Blob storage can include an encryption key on the request for granular control over how blob data is encrypted and decrypted.

Reference: <https://docs.microsoft.com/en-us/azure/storage/common/storage-service-encryption>

QUESTION 14

You plan to deploy an Azure SQL database that will store Personally Identifiable Information (PII). You need to ensure that only privileged users can view the PII.

What should you include in the solution?

- A. Transparent Data Encryption (TDE)
- B. Data Discovery & Classification
- C. dynamic data masking
- D. role-based access control (RBAC)

Answer: C

Explanation:

Dynamic data masking limits sensitive data exposure by masking it to non-privileged users. Dynamic data masking helps prevent unauthorized access to sensitive data by enabling customers to designate how much of the sensitive data to reveal with minimal impact on the application layer. It's a policy-based security feature that hides the sensitive data in the result set of a query over designated database fields, while the data in the database is not changed.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/dynamic-data-masking-overview>

QUESTION 15

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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You plan to deploy multiple instances of an Azure web app across several Azure regions.

You need to design an access solution for the app. The solution must meet the following replication requirements:

- Support rate limiting.
- Balance requests between all instances.
- Ensure that users can access the app in the event of a regional outage.

Solution: You use Azure Load Balancer to provide access to the app.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Azure Application Gateway and Azure Load Balancer do not support rate or connection limits.

Note: Azure Front Door would meet the requirements. The Azure Web Application Firewall (WAF) rate limit rule for Azure Front Door controls the number of requests allowed from clients during a one-minute duration.

Reference:

<https://www.nginx.com/blog/nginx-plus-and-azure-load-balancers-on-microsoft-azure/>
<https://docs.microsoft.com/en-us/azure/web-application-firewall/afds/waf-front-door-rate-limit-powershell>

QUESTION 16

You are planning an Azure IoT Hub solution that will include 50,000 IoT devices. Each device will stream data, including temperature, device ID, and time data. Approximately 50,000 records will be written every second. The data will be visualized in near real time. You need to recommend a service to store and query the data.

Which two services can you recommend? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Azure Table Storage
- B. Azure Event Grid
- C. Azure Cosmos DB SQL API
- D. Azure Time Series Insights

Answer: CD

Explanation:

D: Time Series Insights is a fully managed service for time series data. In this architecture, Time Series Insights performs the roles of stream processing, data store, and analytics and reporting. It accepts streaming data from either IoT Hub or Event Hubs and stores, processes, analyzes, and displays the data in near real time.

C: The processed data is stored in an analytical data store, such as Azure Data Explorer, HBase, Azure Cosmos DB, Azure Data Lake, or Blob Storage.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/scenarios/time-series>

QUESTION 17

You have an on-premises application named App1 that uses an Oracle database. You plan to use Azure Databricks to transform and load data from App1 to an Azure Synapse Analytics instance.

You need to ensure that the App1 data is available to Databricks.

Which two Azure services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Azure Data Box Gateway
- B. Azure Import/Export service
- C. Azure Data Lake Storage
- D. Azure Data Box Edge
- E. Azure Data Factory

Answer: CE

Explanation:

ADF moves data from on-prem Oracle to Data Lake storage, which makes data ready for DataBric

<https://docs.microsoft.com/en-us/azure/data-factory/load-azure-data-lake-storage-gen2>

DataBricks "ETL" data to Synapse:

<https://docs.microsoft.com/en-us/azure/databricks/scenarios/databricks-extract-load-sql-data-warehouse>

QUESTION 18

You have an on-premises line-of-business (LOB) application that uses a Microsoft SQL Server instance as the backend.

You plan to migrate the on-premises SQL Server instance to Azure virtual machines.

You need to recommend a highly available SQL Server deployment that meets the following requirements:

- Minimizes costs
- Minimizes failover time if a single server fails

What should you include in the recommendation?

- A. an Always On availability group that has premium storage disks and a virtual network name (VNN)
- B. an Always On Failover Cluster Instance that has a virtual network name (VNN) and a standard file share
- C. an Always On availability group that has premium storage disks and a distributed network name (DNN)
- D. an Always On Failover Cluster Instance that has a virtual network name (VNN) and a premium file share

Answer: C

Explanation:

Always On availability groups on Azure Virtual Machines are similar to Always On availability groups on-premises, and rely on the underlying Windows Server Failover Cluster.

If you deploy your SQL Server VMs to a single subnet, you can configure a virtual network name

(VNN) and an Azure Load Balancer, or a distributed network name (DNN) to route traffic to your availability group listener.

There are some behavior differences between the functionality of the VNN listener and DNN listener that are important to note:

* Failover time: Failover time is faster when using a DNN listener since there is no need to wait for the network load balancer to detect the failure event and change its routing.

* Etc.

Incorrect:

Not B, not D: Migrate to an Always On availability group, not an Always on Failover cluster Instance.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/availability-group-overview>

QUESTION 19

You need to design a highly available Azure SQL database that meets the following requirements:

- Failover between replicas of the database must occur without any data loss.
- The database must remain available in the event of a zone outage.
- Costs must be minimized.

Which deployment option should you use?

- A. Azure SQL Managed Instance Business Critical
- B. Azure SQL Managed Instance General Purpose
- C. Azure SQL Database Business Critical
- D. Azure SQL Database Serverless

Answer: C

Explanation:

To prevent Data Loss, Premium/Business Critical is required:

The primary node constantly pushes changes to the secondary nodes in order and ensures that the data is persisted to at least one secondary replica before committing each transaction. This process guarantees that if the primary node crashes for any reason, there is always a fully synchronized node to fail over to.

QUESTION 270

You have an Azure web app that uses an Azure key vault named KeyVault1 in the West US Azure region.

You are designing a disaster recovery plan for KeyVault1.

You plan to back up the keys in KeyVault1.

You need to identify to where you can restore the backup.

What should you identify?

- A. any region worldwide
- B. the same region only
- C. KeyVault1 only
- D. the same geography only

Answer: D

Explanation:

When you back up a key vault object, such as a secret, key, or certificate, the backup operation will download the object as an encrypted blob. This blob can't be decrypted outside of Azure. To get usable data from this blob, you must restore the blob into a key vault within the same Azure subscription and Azure geography.

<https://docs.microsoft.com/en-us/azure/key-vault/general/backup?tabs=azure-cli>

QUESTION 21

You have an on-premises application that consumes data from multiple databases. The application code references database tables by using a combination of the server, database, and table name.

You need to migrate the application data to Azure.

To which two services can you migrate the application data to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. SQL Server Stretch Database
- B. SQL Server on an Azure virtual machine
- C. Azure SQL Database
- D. Azure SQL Managed Instance

Answer: BD

Explanation:

Cross-database queries are supported by SQL Server, for example on an Azure virtual machine, and also supported by an Azure SQL Managed Instance.

<https://docs.microsoft.com/en-us/azure/azure-sql/database/elastic-query-getting-started-vertical?view=azuresql>

QUESTION 22

Drag and Drop Question

You are designing a virtual machine that will run Microsoft SQL Server and contain two data disks. The first data disk will store log files, and the second data disk will store data. Both disks are P40 managed disks.

You need to recommend a host caching method for each disk. The method must provide the best overall performance for the virtual machine while preserving the integrity of the SQL data and logs.

Which host caching method should you recommend for each disk? To answer, drag the appropriate methods to the correct disks. Each method 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.

Methods

None
ReadOnly
ReadWrite

Answer Area

Log: Method
Data: Method

Answer:

Methods

ReadWrite

Answer Area

Log: None
Data: ReadOnly

Explanation:

Box 1: None
No data disk caching for the Log files.

Box 2: ReadOnly

Guidelines to optimize performance for your SQL Server on Azure Virtual Machines (VMs) include:

- Set host caching to read-only for data file disks.
- Set host caching to none for log file disks.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/performance-guidelines-best-practices-storage>

Thank You for Trying Our Product

Lead2pass Certification Exam Features:

- ★ More than **99,900** Satisfied Customers Worldwide.
- ★ Average **99.9%** Success Rate.
- ★ **Free Update** to match latest and real exam scenarios.
- ★ **Instant Download** Access! No Setup required.
- ★ Questions & Answers are downloadable in **PDF** format and **VCE** test engine format.
- ★ Multi-Platform capabilities - **Windows, Laptop, Mac, Android, iPhone, iPod, iPad**.
- ★ **100%** Guaranteed Success or **100%** Money Back Guarantee.
- ★ **Fast**, helpful support **24x7**.



View list of all certification exams: <http://www.lead2pass.com/all-products.html>



Microsoft



ORACLE



CITRIX



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14