



**Exam** : **DP-300**

**Title** : Administering Relational  
Databases on Microsoft  
Azure

**Version** : DEMO

## 1. Topic 1, Litware

### **Existing Environment**

#### **Network Environment**

The manufacturing and research datacenters connect to the primary datacenter by using a VPN. The primary datacenter has an ExpressRoute connection that uses both Microsoft peering and private peering. The private peering connects to an Azure virtual network named HubVNet.

#### **Identity Environment**

Litware has a hybrid Azure Active Directory (Azure AD) deployment that uses a domain named litwareinc.com. All Azure subscriptions are associated to the litwareinc.com Azure AD tenant.

#### **Database Environment**

The sales department has the following database workload:

- An on-premises named SERVER1 hosts an instance of Microsoft SQL Server 2012 and two 1-TB databases.
- A logical server named SalesSrv01A contains a geo-replicated Azure SQL database named SalesSQLDb1. SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool. SalesSQLDb1 uses database firewall rules and contained database users.
- An application named SalesSQLDb1App1 uses SalesSQLDb1.

The manufacturing office contains two on-premises SQL Server 2016 servers named SERVER2 and SERVER3. The servers are nodes in the same Always On availability group. The availability group contains a database named ManufacturingSQLDb1

Database administrators have two Azure virtual machines in HubVnet named VM1 and VM2 that run Windows Server 2019 and are used to manage all the Azure databases.

#### **Licensing Agreement**

Litware is a Microsoft Volume Licensing customer that has License Mobility through Software Assurance.

#### **Current Problems**

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent blocking queries.

#### **Requirements**

##### **Planned Changes**

Litware plans to implement the following changes:

- Implement 30 new databases in Azure, which will be used by time-sensitive manufacturing apps that have varying usage patterns. Each database will be approximately 20 GB.
- Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01. ResearchDB1 will contain Personally Identifiable Information (PII) data.
- Develop an app named ResearchApp1 that will be used by the research department to populate and access ResearchDB1.

- Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.
- Migrate the SERVER1 databases to the Azure SQL Database platform.

### **Technical Requirements**

Litware identifies the following technical requirements:

- Maintenance tasks must be automated.
- The 30 new databases must scale automatically.
- The use of an on-premises infrastructure must be minimized.
- Azure Hybrid Use Benefits must be leveraged for Azure SQL Database deployments.
- All SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

### **Security and Compliance Requirements**

Litware identifies the following security and compliance requirements:

- Store encryption keys in Azure Key Vault.
- Retain backups of the PII data for two months.
- Encrypt the PII data at rest, in transit, and in use.
- Use the principle of least privilege whenever possible.
- Authenticate database users by using Active Directory credentials.
- Protect Azure SQL Database instances by using database-level firewall rules.
- Ensure that all databases hosted in Azure are accessible from VM1 and VM2 without relying on public endpoints.

### **Business Requirements**

Litware identifies the following business requirements:

- Meet an SLA of 99.99% availability for all Azure deployments.
- Minimize downtime during the migration of the SERVER1 databases.
- Use the Azure Hybrid Use Benefits when migrating workloads to Azure.
- Once all requirements are met, minimize costs whenever possible.

### **HOTSPOT**

You are planning the migration of the SERVER1 databases. The solution must meet the business requirements.

What should you include in the migration plan? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

## Answer Area

Azure Database Migration

Service pricing tier:

	▼
Standard 2-vCore	
Standard 4-vCore	
Premium 4-vCore	

Required Azure resource:

	▼
A virtual network that has service endpoints	
A VPN gateway	
An Azure Logic app	

Answer:

## Answer Area

Azure Database Migration

Service pricing tier:

	▼
Standard 2-vCore	
Standard 4-vCore	
Premium 4-vCore	

Required Azure resource:

	▼
A virtual network that has service endpoints	
A VPN gateway	
An Azure Logic app	

### Explanation:

Azure Database Migration service

Box 1: Premium 4-VCore

Scenario: Migrate the SERVER1 databases to the Azure SQL Database platform.

⇒ Minimize downtime during the migration of the SERVER1 databases.

Premium 4-vCore is for large or business critical workloads. It supports online migrations, offline migrations, and faster migration speeds.

Incorrect Answers:

The Standard pricing tier suits most small- to medium- business workloads, but it supports offline migration only.

Box 2: A VPN gateway

You need to create a Microsoft Azure Virtual Network for the Azure Database Migration Service by using the Azure Resource Manager deployment model, which provides site-to-site connectivity to your on-premises source servers by using either ExpressRoute or VPN.

Reference:

<https://azure.microsoft.com/pricing/details/database-migration/>

<https://docs.microsoft.com/en-us/azure/dms/tutorial-sql-server-azure-sql-online>

## 2.DRAG DROP

You need to configure user authentication for the SERVER1 databases. The solution must meet the security and compliance requirements.

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.

### Actions

Create a user in the master database

Modify the Azure SQL server administrator account

Create contained database users

Create an Azure AD administrator for the logical server

Connect to the databases by using an Azure AD account

Enable the contained database authentication option

### Answer Area

Answer:

### Actions

Create a user in the master database

Modify the Azure SQL server administrator account

Create contained database users

Create an Azure AD administrator for the logical server

Connect to the databases by using an Azure AD account

Enable the contained database authentication option

### Answer Area

Create an Azure AD administrator for the logical server

Create contained database users

Connect to the databases by using an Azure AD account

**Explanation:**

Scenario: Authenticate database users by using Active Directory credentials.

The configuration steps include the following procedures to configure and use Azure Active Directory authentication.

Create and populate Azure AD.

Optional: Associate or change the active directory that is currently associated with your Azure Subscription.

Create an Azure Active Directory administrator. (Step 1)

Configure your client computers.

Create contained database users in your database mapped to Azure AD identities. (Step 2)

Connect to your database by using Azure AD identities. (Step 3)

Reference: <https://docs.microsoft.com/en-us/azure/azure-sql/database/authentication-aad-overview>

**3.HOTSPOT**

You need to implement the monitoring of SalesSQLDb1. The solution must meet the technical requirements.

How should you collect and stream metrics? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

**Answer Area**

Collect metrics from:

	▼
The database only	
The elastic pool and the database	
The elastic pool only	
The server, the elastic pool, and the database	

Stream metrics to:

	▼
Azure Event Hubs	
Azure Log Analytics	
Azure Storage	

**Answer:**

## Answer Area

Collect metrics from:

	▼
The database only	
The elastic pool and the database	
The elastic pool only	
The server, the elastic pool, and the database	

Stream metrics to:

	▼
Azure Event Hubs	
Azure Log Analytics	
Azure Storage	

### Explanation:

Box 1: The server, the elastic pool, and the database

Scenario:

SalesSQLDb1 is in an elastic pool named SalesSQLDb1Pool.

Litware technical requirements include: all SQL Server and Azure SQL Database metrics related to CPU and storage usage and limits must be analyzed by using Azure built-in functionality.

Box 2: Azure Event hubs

Scenario: Migrate ManufacturingSQLDb1 to the Azure virtual machine platform.

Event hubs are able to handle custom metrics.

Incorrect Answers:

Azure Log Analytics

Azure metric and log data are sent to Azure Monitor Logs, previously known as Azure Log Analytics, directly by Azure. Azure SQL Analytics is a cloud only monitoring solution supporting streaming of diagnostics telemetry for all of your Azure SQL databases.

However, because Azure SQL Analytics does not use agents to connect to Azure Monitor, it does not support monitoring of SQL Server hosted on-premises or in virtual machines.

4.You need to identify the cause of the performance issues on SalesSQLDb1.

Which two dynamic management views should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. sys.dm\_pdw\_nodes\_tran\_locks
- B. sys.dm\_exec\_compute\_node\_errors
- C. sys.dm\_exec\_requests
- D. sys.dm\_cdc\_errors
- E. sys.dm\_pdw\_nodes\_os\_wait\_stats
- F. sys.dm\_tran\_locks

**Answer:** AE

### Explanation:

SalesSQLDb1 experiences performance issues that are likely due to out-of-date statistics and frequent

blocking queries.

A: Use `sys.dm_pdw_nodes_tran_locks` instead of `sys.dm_tran_locks` from Azure Synapse Analytics (SQL Data Warehouse) or Parallel Data Warehouse.

E: Example:

The following query will show blocking information.

```
SELECT
t1.resource_type,
t1.resource_database_id,
t1.resource_associated_entity_id,
t1.request_mode,
t1.request_session_id,
t2.blocking_session_id
FROM sys.dm_tran_locks as t1
INNER JOIN sys.dm_os_waiting_tasks as t2
ON t1.lock_owner_address = t2.resource_address;
```

Note: Depending on the system you're working with you can access these wait statistics from one of three locations:

`sys.dm_os_wait_stats`: for SQL Server

`sys.dm_db_wait_stats`: for Azure SQL Database

`sys.dm_pdw_nodes_os_wait_stats`: for Azure SQL Data Warehouse

Incorrect Answers:

F: `sys.dm_tran_locks` returns information about currently active lock manager resources in SQL Server 2019 (15.x). Each row represents a currently active request to the lock manager for a lock that has been granted or is waiting to be granted.

Instead use `sys.dm_pdw_nodes_tran_locks`.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-tran-locks-transact-sql>

## 5.HOTSPOT

You need to recommend a configuration for ManufacturingSQLDb1 after the migration to Azure. The solution must meet the business requirements.

What should you include in the recommendation? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

## Answer Area

Quorum model:

	▼
Cloud witness	
Disk witness	
File share witness	

Azure resource for the availability group listener:

	▼
Azure Application Gateway	
Azure Basic Load Balancer	

Answer:

## Answer Area

Quorum model:

	▼
Cloud witness	
Disk witness	
File share witness	

Azure resource for the availability group listener:

	▼
Azure Application Gateway	
Azure Basic Load Balancer	

### Explanation:

Scenario: Business Requirements

Litware identifies business requirements include: meet an SLA of 99.99% availability for all Azure deployments.

Box 1: Cloud witness

If you have a Failover Cluster deployment, where all nodes can reach the internet (by extension of Azure), it is recommended that you configure a Cloud Witness as your quorum witness resource.

Box 2: Azure Basic Load Balancer

Microsoft guarantees that a Load Balanced Endpoint using Azure Standard Load Balancer, serving two or more Healthy Virtual Machine Instances, will be available 99.99% of the time.

Note: There are two main options for setting up your listener: external (public) or internal. The external (public) listener uses an internet facing load balancer and is associated with a public Virtual IP (VIP) that is accessible over the internet. An internal listener uses an internal load balancer and only supports clients within the same Virtual Network.

Reference: <https://technet.microsoft.com/windows-server-docs/failover-clustering/deploy-cloud-witness>

[https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1\\_0/](https://azure.microsoft.com/en-us/support/legal/sla/load-balancer/v1_0/)