



Vendor: Microsoft

Exam Code: DP-300

Exam Name: Administering Relational Databases on
Microsoft Azure

Version: DEMO

QUESTION 1

Case Study 1 - Litware, Inc

Overview

Litware, Inc. is a renewable energy company that has a main office in Boston. The main office hosts a sales department and the primary datacenter for the company.

Physical Locations

Existing Environment

Litware has a manufacturing office and a research office in separate locations near Boston. Each office has its own datacenter and internet connection.

The manufacturing and research datacenters connect to the primary datacenter by using a VPN.

Network Environment

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.

You need to implement authentication for ResearchDB1. The solution must meet the security and compliance requirements.

What should you run as part of the implementation?

- A. CREATE LOGIN and the FROM WINDOWS clause
- B. CREATE USER and the FROM CERTIFICATE clause
- C. CREATE USER and the FROM LOGIN clause
- D. CREATE USER and the ASYMMETRIC KEY clause
- E. CREATE USER and the FROM EXTERNAL PROVIDER clause

Answer: E

Explanation:

Scenario: Authenticate database users by using Active Directory credentials.

(Create a new Azure SQL database named ResearchDB1 on a logical server named ResearchSrv01.)

Authenticate the user in SQL Database or SQL Data Warehouse based on an Azure Active Directory user:

```
CREATE USER [Fritz@contoso.com] FROM EXTERNAL PROVIDER;
```

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-user-transact-sql>

QUESTION 2

Case Study 2 - Contoso, Ltd

Overview

General Overview

Contoso, Ltd. is a financial data company that has 100 employees. The company delivers financial data to customers.

Physical Locations

Contoso has a datacenter in Los Angeles and an Azure subscription. All Azure resources are in the US West 2 Azure region. Contoso has a 10-Gb ExpressRoute connection to Azure.

The company has customers worldwide.

Existing Environment

Active Directory

Contoso has a hybrid Azure Active Directory (Azure AD) deployment that syncs to on-premises Active Directory.

Database Environment

Contoso has SQL Server 2017 on Azure virtual machines shown in the following table.

Name	Role
SQL1	Primary data warehouse
SQL2	Secondary data warehouse
SQL3	Extract, transform, and load (ETL) server

SQL1 and SQL2 are in an Always On availability group and are actively queried. SQL3 runs jobs, provides historical data, and handles the delivery of data to customers.

What should you implement to meet the disaster recovery requirements for the PaaS solution?

- A. Availability Zones
- B. failover groups
- C. Always On availability groups
- D. geo-replication

Answer: B

Explanation:

Scenario: In the event of an Azure regional outage, ensure that the customers can access the PaaS solution with minimal downtime. The solution must provide automatic failover.

The auto-failover groups feature allows you to manage the replication and failover of a group of databases on a server or all databases in a managed instance to another region. It is a declarative abstraction on top of the existing active geo-replication feature, designed to simplify deployment and management of geo-replicated databases at scale. You can initiate failover manually or you can delegate it to the Azure service based on a user-defined policy.

The latter option allows you to automatically recover multiple related databases in a secondary region after a catastrophic failure or other unplanned event that results in full or partial loss of the SQL Database or SQL Managed Instance availability in the primary region.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/auto-failover-group-overview>

QUESTION 3

Case Study 3 - Contoso, Ltd 2

Overview

Contoso, Ltd. is a clothing retailer based in Seattle. The company has 2,000 retail stores across the United States and an emerging online presence.

The network contains an Active Directory forest named contoso.com. The forest is integrated with an Azure Active Directory (Azure AD) tenant named contoso.com. Contoso has an Azure subscription associated to the contoso.com Azure AD tenant.

Existing Environment

Transactional Data

Contoso has three years of customer, transaction, operational, sourcing, and supplier data comprised of 10 billion records stored across multiple on-premises Microsoft SQL Server servers. The SQL Server instances contain data from various operations systems. The data is loaded into the instances by using SQL Server Integration Services (SSIS) packages.

You estimate that combining all product sales transactions into a company-wide sales transactions dataset will result in a single table that contains 5 billion rows, with one row per transaction.

Most queries targeting the sales transactions data will be used to identify which products were sold in retail stores and which products were sold online during different time periods. Sales transaction data that is older than three years will be removed monthly.

You plan to create a retail store table that will contain the address of each retail store. The table will be approximately 2 MB. Queries for retail store sales will include the retail store addresses.

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

You need to design a data retention solution for the Twitter feed data records. The solution must meet the customer sentiment analytics requirements.

Which Azure Storage functionality should you include in the solution?

- A. time-based retention
- B. change feed
- C. lifecycle management
- D. soft delete

Answer: C

Explanation:

The lifecycle management policy lets you:

- Delete blobs, blob versions, and blob snapshots at the end of their lifecycles

Scenario:

- Purge Twitter feed data records that are older than two years.
- Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.
- Minimize administrative effort to maintain the Twitter feed data records.

Incorrect Answers:

A: Time-based retention policy support: Users can set policies to store data for a specified interval. When a time-based retention policy is set, blobs can be created and read, but not modified or deleted. After the retention period has expired, blobs can be deleted but not overwritten.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts>

QUESTION 4

Case Study 4 - A.Datum

Overview

ADatum Corporation is a retailer that sells products through two sales channels: retail stores and a website.

Existing Environment

ADatum has one database server that has Microsoft SQL Server 2016 installed. The server hosts three mission-critical databases named SALESDB, DOCDB, and REPORTINGDB.

SALESDB collects data from the stores and the website.

DOCDB stores documents that connect to the sales data in SALESDB. The documents are stored in two different JSON formats based on the sales channel.

REPORTINGDB stores reporting data and contains several columnstore indexes. A daily process creates reporting data in REPORTINGDB from the data in SALESDB. The process is implemented as a SQL Server Integration Services (SSIS) package that runs a stored procedure from SALESDB.

Requirements

Planned Changes

ADatum plans to move the current data infrastructure to Azure. The new infrastructure has the following requirements:

- Migrate SALESDB and REPORTINGDB to an Azure SQL database.
- Migrate DOCDB to Azure Cosmos DB.
- The sales data, including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics. The analytics process will perform aggregations that must be done continuously, without gaps, and without overlapping.
- As they arrive, all the sales documents in JSON format must be transformed into one consistent format.
- Azure Data Factory will replace the SSIS process of copying the data from SALESDB to REPORTINGDB.

Which counter should you monitor for real-time processing to meet the technical requirements?

- A. SU% Utilization
- B. CPU% utilization
- C. Concurrent users
- D. Data Conversion Errors

Answer: A

Explanation:

Scenario:

- Real-time processing must be monitored to ensure that workloads are sized properly based on actual usage patterns.

- The sales data including the documents in JSON format, must be gathered as it arrives and analyzed online by using Azure Stream Analytics.

Streaming Units (SUs) represents the computing resources that are allocated to execute a Stream Analytics job. The higher the number of SUs, the more CPU and memory resources are allocated for your job. This capacity lets you focus on the query logic and abstracts the need to manage the hardware to run your Stream Analytics job in a timely manner.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-streaming-unit-consumption>

QUESTION 5

You have an Azure virtual machine named VM1 on a virtual network named VNet1. Outbound traffic from VM1 to the internet is blocked.

You have an Azure SQL database named SqlDb1 on a logical server named SqlSrv1.

You need to implement connectivity between VM1 and SqlDb1 to meet the following requirements:

- Ensure that VM1 cannot connect to any Azure SQL Server other than SqlSrv1.
- Restrict network connectivity to SqlSrv1.

What should you create on VNet1?

- A. a VPN gateway
- B. a service endpoint
- C. a private endpoint
- D. an ExpressRoute gateway

Answer: C

Explanation:

A private endpoint is a network interface that uses a private IP address from your virtual network. This network interface connects you privately and securely to a service powered by Azure Private Link. By enabling a private endpoint, you're bringing the service into your virtual network.

The service could be an Azure service such as:

- Azure Storage
- Azure Cosmos DB
- Azure SQL Database
- Your own service using a Private Link Service.

<https://docs.microsoft.com/en-us/azure/private-link/private-endpoint-overview>

QUESTION 6

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 have SQL Server 2019 on an Azure virtual machine.

You are troubleshooting performance issues for a query in a SQL Server instance.

To gather more information, you query sys.dm_exec_requests and discover that the wait type is PAGELATCH_UP and the wait_resource is 2:3:905856.

You need to improve system performance.

Solution: You create additional tempdb files.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

To improve the concurrency of tempdb, try the following methods:

- * Increase the number of data files in tempdb to maximize disk bandwidth and reduce contention in allocation structures.
- * Etc.

Note: Symptoms On a server that is running Microsoft SQL Server, you notice severe blocking when the server is experiencing a heavy load. Dynamic Management Views [sys.dm_exec_request or sys.dm_os_waiting_tasks] indicates that these requests or tasks are waiting for tempdb resources. Additionally, the wait type is PAGELATCH_UP, and the wait resource points to pages in Tempdb.

Reference:

<https://docs.microsoft.com/en-US/troubleshoot/sql/performance/recommendations-reduce-allocation-contention>

QUESTION 7

Hotspot Question

You have an Azure Data Lake Storage Gen2 account named account1 that stores logs as shown in the following table.

Type	Designated retention period
Application	360 days
Infrastructure	60 days

You do not expect that the logs will be accessed during the retention periods.

You need to recommend a solution for account1 that meets the following requirements:

- Automatically deletes the logs at the end of each retention period
- Minimizes storage costs

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

To minimize storage costs:

Store the infrastructure logs and the application logs in the Archive access tier.
Store the infrastructure logs and the application logs in the Cool access tier.
Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

Azure Data Factory pipelines
Azure Blob storage lifecycle management rules
Immutable Azure Blob storage time-based retention policies

Answer:

Answer Area

To minimize storage costs:

Store the infrastructure logs and the application logs in the Archive access tier.
Store the infrastructure logs and the application logs in the Cool access tier.
Store the infrastructure logs in the Cool access tier and the application logs in the Archive access tier.

To delete the logs automatically:

Azure Data Factory pipelines
Azure Blob storage lifecycle management rules
Immutable Azure Blob storage time-based retention policies

Explanation:

Box 1: Store the infrastructure logs in the Cool access tier the application logs in the Archive access tier

Hot - Optimized for storing data that is accessed frequently.

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

Box 2: Azure Blob storage lifecycle management rules

Blob storage lifecycle management offers a rich, rule-based policy that you can use to transition your data to the best access tier and to expire data at the end of its lifecycle.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

QUESTION 8

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 have an Azure Data Lake Storage account that contains a staging zone.

You need to design a daily process to ingest incremental data from the staging zone, transform the data by executing an R script, and then insert the transformed data into a data warehouse in Azure Synapse Analytics.

Solution: You use an Azure Data Factory schedule trigger to execute a pipeline that copies the data to a staging table in the data warehouse, and then uses a stored procedure to execute the R script.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Azure Synapse Analytics does not support R script.

Correct solution: You use an Azure Data Factory schedule trigger to execute a pipeline that executes an Azure Databricks notebook, and then inserts the data into the data warehouse.

Reference:

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

QUESTION 9

You have an Azure Data Factory that contains 10 pipelines.

You need to label each pipeline with its main purpose of either ingest, transform, or load. The labels must be available for grouping and filtering when using the monitoring experience in Data Factory.

What should you add to each pipeline?

- A. an annotation
- B. a resource tag
- C. a run group ID
- D. a user property
- E. a correlation ID

Answer: A

Explanation:

Azure Data Factory annotations help you easily filter different Azure Data Factory objects based on a tag.

You can define tags so you can see their performance or find errors faster.

Reference:

<https://www.techtalkcorner.com/monitor-azure-data-factory-annotations/>

QUESTION 10

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies queries that cause performance issues due to tempDB contention.

You need to resolve the performance issues.

What should you do?

- A. Implement memory-optimized tables.
- B. Run the dbcc flushprocindb command.
- C. Replace the sequential index keys with nonsequential keys.
- D. Run the dbcc dbreindex command.

Answer: A

Explanation:

TempDB contention troubleshooting:

The diagnostics log outputs tempDB contention details. You can use the information as the starting point for troubleshooting. There are two things you can pursue to alleviate this kind of contention and increase the throughput of the overall workload: You can stop using the temporary tables. You also can use memory- optimized tables.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/intelligent-insights-troubleshoot-performance#tempdb-contention>

QUESTION 11

You have an Azure SQL managed instance.

You need to enable SQL Agent Job email notifications.

What should you do?

- A. Use the Agent XPs option.
- B. Enable the SQL Server Agent.
- C. Run the sp_configure command.
- D. Run the sp_set_agent_properties command.

Answer: C

Explanation:

You would need to enable Database email extended procedure using Database Mail XPs configuration option:

EXEC sp_configure 'show advanced options', 1;

GO -

RECONFIGURE;

GO -

EXEC sp_configure 'Database Mail XPs', 1;

GO -

RECONFIGURE -

GO -

Now you can test the configuration by sending emails using sp_send and sp_notify_operator procedures.

Reference:

<https://techcommunity.microsoft.com/t5/azure-sql-blog/sending-emails-in-azure-sql-managed-instance/ba-p/386235>

QUESTION 12

You have SQL Server on an Azure virtual machine.

You need to add a 4-TB volume that meets the following requirements:

- Maximizes IOPs
- Uses premium solid state drives (SSDs)

What should you do?

- A. Attach two mirrored 4-TB SSDs.
- B. Attach a stripe set that contains four 1-TB SSDs.
- C. Attach a RAID-5 array that contains five 1-TB SSDs.
- D. Attach a single 4-TB SSD.

Answer: B

Explanation:

For more throughput, you can add additional data disks and use disk striping.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/virtual-machines/windows/storage-configuration?tabs=windows2016>

QUESTION 13

You have an Azure SQL database named db1 on a server named server1.

The Intelligent Insights diagnostics log identifies that several tables are missing indexes.

You need to ensure that indexes are created for the tables.

What should you do?

- A. Run the DBCC SQLPERF command.
- B. Run the DBCC DBREINDEX command.
- C. Modify the automatic tuning settings for db1.
- D. Modify the Query Store settings for db1.

Answer: C

Explanation:

Automatic tuning is a fully managed intelligent performance service that uses built-in intelligence to continuously monitor queries executed on a database, and it automatically improves their performance.

Automatic tuning for Azure SQL Database uses the CREATE INDEX, DROP INDEX, and FORCE LAST GOOD PLAN database advisor recommendations to optimize your database performance.

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/automatic-tuning-overview>

QUESTION 14

Hotspot Question

You have a SQL Server on Azure Virtual Machines instance that hosts a 10-TB SQL database named DB1.

You need to identify and repair any physical or logical corruption in DB1. The solution must meet the following requirements:

- Minimize how long it takes to complete the procedure.
- Minimize data loss.

How should you complete the command? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

DBCC CHECK [DB1],

▼
NOINDEX
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
REPAIR_REBUILD

) WITH

▼
EXTENDED_LOGICAL_CHECKS;
PHYSICAL_ONLY;
TABLOCK;

Answer:

Answer Area

DBCC CHECK [DB1],

▼
NOINDEX
REPAIR_ALLOW_DATA_LOSS
REPAIR_FAST
REPAIR_REBUILD

) WITH

▼
EXTENDED_LOGICAL_CHECKS;
PHYSICAL_ONLY;
TABLOCK;

Explanation:

Box 1: REPAIR_REBUILD

Performs repairs that have no possibility of data loss. This option may include quick repairs, such as repairing missing rows in nonclustered indexes, and more time-consuming repairs, such as rebuilding an index.

Box 2: PHYSICAL_ONLY

Limits the checking to the integrity of the physical structure of the page and record headers and the allocation consistency of the database. This check is designed to provide a small overhead check of the physical consistency of the database, but it can also detect torn pages, checksum failures, and common hardware failures that can compromise a user's data.

Incorrect:

TABLOCK

Causes DBCC CHECKDB to obtain locks instead of using an internal database snapshot. This includes a short-term exclusive (X) lock on the database.

TABLOCK will cause DBCC CHECKDB to run faster on a database under heavy load, but will decrease the concurrency available on the database while DBCC CHECKDB is running.

EXTENDED_LOGICAL_CHECKS

If the compatibility level is 100 (SQL Server 2008) or higher, performs logical consistency checks on an indexed view, XML indexes, and spatial indexes, where present.

Reference:

<https://docs.microsoft.com/en-us/sql/t-sql/database-console-commands/dbcc-checkdb-transact-sql>

QUESTION 15

Your company uses Azure Stream Analytics to monitor devices.

The company plans to double the number of devices that are monitored.

You need to monitor a Stream Analytics job to ensure that there are enough processing resources to handle the additional load.

Which metric should you monitor?

- A. Input Deserialization Errors
- B. Late Input Events
- C. Early Input Events
- D. Watermark delay

Answer: D

Explanation:

The Watermark delay metric is computed as the wall clock time of the processing node minus the largest watermark it has seen so far.

The watermark delay metric can rise due to:

1. Not enough processing resources in Stream Analytics to handle the volume of input events.
2. Not enough throughput within the input event brokers, so they are throttled.
3. Output sinks are not provisioned with enough capacity, so they are throttled.

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-time-handling>

QUESTION 16

You manage an enterprise data warehouse in Azure Synapse Analytics.

Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries.

You need to monitor resource utilization to determine the source of the performance issues.

Which metric should you monitor?

- A. Local tempdb percentage
- B. DWU percentage
- C. Data Warehouse Units (DWU) used
- D. Cache hit percentage

Answer: D

Explanation:

You can use Azure Monitor to view cache metrics to troubleshoot query performance.

The key metrics for troubleshooting the cache are Cache hit percentage and Cache used percentage.

Possible scenario: Your current working data set cannot fit into the cache which causes a low cache hit percentage due to physical reads. Consider scaling up your performance level and rerun your workload to populate the cache.

Reference:

<https://docs.microsoft.com/da-dk/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-how-to-monitor-cache>

QUESTION 17

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 and a database named DB1.

DB1 contains a fact table named Table.

You need to identify the extent of the data skew in Table1.

What should you do in Synapse Studio?

- A. Connect to Pool1 and query sys.dm_pdw_nodes_db_partition_stats.
- B. Connect to the built-in pool and run DBCC CHECKALLOC.
- C. Connect to Pool1 and run DBCC CHECKALLOC.
- D. Connect to the built-in pool and query sys.dm_pdw_nodes_db_partition_stats.

Answer: A

Explanation:

First connect to Pool1, not the built-in serverless pool, then use sys.dm_pdw_nodes_db_partition_stats to analyze any skewness in the data.

Reference:

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/cheat-sheet>

QUESTION 18

You have an Azure subscription that contains a logical SQL server named Server1.

The master database of Server1 contains a user named User1.

You need to ensure that User1 can create databases on Server1.

Which database role should you assign to User1?

- A. db_owner
- B. dbmanager
- C. dbo
- D. db_ddladmin

Answer: B

Explanation:

dbmanager: Can create and delete databases. A member of the dbmanager role that creates a database, becomes the owner of that database, which allows that user to connect to that database as the dbo user. The dbo user has all database permissions in the database. Members of the dbmanager role don't necessarily have permission to access databases that they don't own.

Reference:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/authentication-access/database-level-roles>

QUESTION 19

You have a on-premises Microsoft SQL Server named SQL1 that hosts five databases.

You need to migrate the databases to an Azure SQL managed instance.

The solution must minimize downtime and prevent data loss.

What should you use?

- A. log shipping
- B. Always On availability groups
- C. Database Migration Assistant
- D. Backup and Restore

Answer: C

Explanation:

The Data Migration Assistant (DMA) helps you upgrade to a modern data platform by detecting compatibility issues that can impact database functionality in your new version of SQL Server or Azure SQL Database. DMA recommends performance and reliability improvements for your target environment and allows you to move your schema, data, and uncontained objects from your source server to your target server.

Capabilities include:

Assess on-premises SQL Server instance(s) migrating to Azure SQL database(s).

Note: For large migrations (in terms of number and size of databases), we recommend that you use the Azure Database Migration Service, which can migrate databases at scale.

Migrate an on-premises SQL Server instance to a modern SQL Server instance hosted on-premises or on an Azure virtual machine (VM) that is accessible from your on-premises network.

Reference:

<https://docs.microsoft.com/en-us/sql/dma/dma-overview>

<https://docs.microsoft.com/en-us/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-guide>

QUESTION 20

You have an on-premises Microsoft SQL Server 2019 instance named SQL1 that hosts a database named db1.

You have an Azure subscription that contains an

Azure SQL managed instance named MI1 and an Azure Storage account named storage1.

You plan to migrate db1 to MI1 by using the backup and restore process.

You need to ensure that you can back up db1 to storage1. The solution must meet the following requirements:

- Use block blob storage.
- Maximize security.

What should you do on storage1?

- A. Generate a shared access signature (SAS).
- B. Create an access policy.
- C. Rotate the storage keys.
- D. Enable infrastructure encryption.

Answer: A

Explanation:

For backup and restore process first you need to backup to Azure Storage and in order to authenticate to Azure Storage first you need to create SAS through which authentication would be done and .bak file would be created on Azure storage.

Reference:

<https://learn.microsoft.com/en-us/azure/azure-sql/migration-guides/managed-instance/sql-server-to-managed-instance-guide?view=azuresql#backup-and-restore>

Thank You for Trying Our Product

Passleader 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.passleader.com/all-products.html>



Microsoft



ORACLE



JUNIPER
NETWORKS



EMC²
where information lives

10% Discount Coupon Code: ASTR14