

OFFICIAL MICROSOFT LEARNING PRODUCT

# 20465D

**Designing a Data Solution with Microsoft  
SQL Server 2014**

Information in this document, including URL and other Internet Web site references, is subject to change without notice. Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places, and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, e-mail address, logo, person, place or event is intended or should be inferred. Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

The names of manufacturers, products, or URLs are provided for informational purposes only and Microsoft makes no representations and warranties, either expressed, implied, or statutory, regarding these manufacturers or the use of the products with any Microsoft technologies. The inclusion of a manufacturer or product does not imply endorsement of Microsoft of the manufacturer or product. Links may be provided to third party sites. Such sites are not under the control of Microsoft and Microsoft is not responsible for the contents of any linked site or any link contained in a linked site, or any changes or updates to such sites. Microsoft is not responsible for webcasting or any other form of transmission received from any linked site. Microsoft is providing these links to you only as a convenience, and the inclusion of any link does not imply endorsement of Microsoft of the site or the products contained therein.

© 2016 Microsoft Corporation. All rights reserved.

Microsoft and the trademarks listed at <http://www.microsoft.com/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other trademarks are property of their respective owners

Product Number: 20465D

Part Number (if applicable):

Released: 03/2016

# Module 1

## Introduction to Enterprise Data Architecture

### Contents:

Lesson 1: Considerations for Enterprise Data	2
Lesson 2: Assessing an Existing Infrastructure	4
Module Review and Takeaways	7
Lab Review Questions and Answers	8

## Lesson 1

# Considerations for Enterprise Data

### Contents:

Question and Answers

3

## Question and Answers

**Question:** Which of the challenges highlighted in this lesson have you come across in your own organization? Are there any enterprise management challenges you have experienced that have not been listed?

**Answer:** Answers will vary.

## Lesson 2

# Assessing an Existing Infrastructure

### Contents:

Question and Answers	5
Demonstration: Using the MAP Toolkit	5

## Question and Answers

**Question:** Which of the following is not a technique that the MAP tool can use to discover servers?

- ( ) Windows networking protocols.
- ( ) System Center Configuration Manager (SCCM)
- ( ) Internet Printing Protocol
- ( ) Active Directory Domain Services
- ( ) IP address range scanning

**Answer:**

- ( ) Windows networking protocols.
- ( ) System Center Configuration Manager (SCCM)
- (v) Internet Printing Protocol
- ( ) Active Directory Domain Services
- ( ) IP address range scanning

## Demonstration: Using the MAP Toolkit

### Demonstration Steps

Install the MAP Toolkit

1. Ensure that the MSL-TMG, 20465C-MIA-DC, and 20465C-MIA-SQL virtual machines are running, and log on to 20465C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. Start Internet Explorer and browse to **www.microsoft.com/map**.
3. Click the link to download the latest version of the MAP Toolkit.
4. Follow the instructions to download and install the MAP Toolkit on the MIA-SQL server.
5. On the Start screen, type **MAP** and then click **Microsoft Assessment and Planning Toolkit**.
6. In the **Microsoft Assessment and Planning Toolkit** dialog box, in the **Create or select a database** area, in the **Name** field, type **MAPData**, and then click **OK**.

Discover Database Servers

1. On the **Database** Tab, click **Collect inventory data**.
2. In the **Inventory and Assessment Wizard** dialog box, on the **Inventory Scenarios** page, select **SQL Server with Database Details**. Then click **Next**.
3. On the **Discovery Methods** page, ensure that only **Use Active Directory Domain Services (AD DS)** is selected, and click **Next**.
4. On the **Active Directory Credentials** page, enter the following details and click **Next**:
  - o **Domain:** adventureworks.msft
  - o **Domain account:** ADVENTUREWORKS\Student
  - o **Password:** Pa\$\$w0rd
5. On the **Active Directory Options** page, ensure that **Find all computers in all domains, containers, and organizational units** is selected, and click **Next**.

6. On the **All Computers Credentials** page, click **Create**. Add the following account for WMI and SQL windows technologies and click **Save**. Then click **Next**:
  - **Account Name:** ADVENTUREWORKS\Student
  - **Password:** Pa\$\$w0rd
7. On the **Credentials Order** page, click **Next**.
8. On the **Summary** page, click **Finish** and wait for data collection to complete. Then click **Close**.

#### Generate Reports

1. In the MAP Toolkit, on the **Database** tab, view the information in the **SQL Server Products** tile.
2. Click the **SQL Server Products** tile, and view the summary details that are displayed.
3. In the **Options** area, click **Generate SQL Server Database Details Reports**.
4. When the reports have been generated, click **Close** and view the contents of the reports folder.
5. Open each of the reports in Microsoft Excel and view the details they contain.



## Module Review and Takeaways

**Question:** Based on your own experience, what challenges face a database administrator or data architect in a large enterprise?

**Answer:** Answers will vary depending on the experience of the students.

# Lab Review Questions and Answers

## Lab: Assessing an Existing Enterprise Data Infrastructure

### Question and Answers

#### Lab Review

**Question:** How might the information in the reports generated by the MAP Toolkit be useful to an enterprise data architect?

**Answer:** Answers will vary depending on the experience of the students.

# Module 2

## Multi-Server Configuration Management

### Contents:

Lesson 1: Policy-Based Management	2
Lesson 2: Microsoft System Center	5
Module Review and Takeaways	7

## Lesson 1

# Policy-Based Management

### Contents:

Question and Answers	3
Resources	3
Demonstration: Viewing Best-practice Policies	3

## Question and Answers

**Question:** You have created several policies in a new policy category called "Required". You want to ensure that these policies apply to all your databases. What option must you use?

- ( ) Use the Integrated condition
- ( ) Use the On Change: Prevent evaluation mode
- ( ) Use the On Demand evaluation mode
- ( ) Use the Mandate option.
- ( ) Use the syspolicy\_policy\_execution\_history system stored procedure.

**Answer:**

- ( ) Use the Integrated condition
- ( ) Use the On Change: Prevent evaluation mode
- ( ) Use the On Demand evaluation mode
- (v) Use the Mandate option.
- ( ) Use the syspolicy\_policy\_execution\_history system stored procedure.

## Resources

### Policy-Based Management Concepts



**Reference Links:** For more information about configuring alerts for policy violations, see the topic *Configure Alerts to Notify Policy Administrators of Policy Failures* in SQL Server Books Online.

### Planning a Policy-Based Management Strategy



**Reference Links:** For more information about Policy-Based Management system views, see the *Policy-Based Management Views (Transact-SQL)* topic in SQL Server Books Online.

## Demonstration: Viewing Best-practice Policies

### Demonstration Steps

View Best-Practice Policies

1. Ensure that the 20465C-MIA-DC and 20465C-MIA-SQL virtual machines are running, and then log on to 20465C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. On the taskbar, click **SQL Server Management Studio**.
3. In the **Connect to Server** dialog box, in the **Server name** box, type **MIA-SQL**, and then click **Connect**.
4. In Object Explorer, expand **Management**, expand **Policy Management**, right-click **Policies**, and then click **Import Policy**.
5. In the **Import** dialog box, next to the **Files to import** box, click the ellipsis button.

6. In the **Select Policy** dialog box, double-click **SQL Server Best Practices**, double-click **DatabaseEngine**, and then double-click **1033**.
7. Browse the list of best-practice policies.
8. In the **Select Policy** dialog box, click **Cancel**.
9. In the **Import** dialog box, click **Cancel**.

## Lesson 2

# Microsoft System Center

### Contents:

Question and Answers

6

## Question and Answers

**Question:** Your organization uses Microsoft System Center to manage its Windows servers and desktop infrastructure. You have been asked to begin rationalizing the organization's database infrastructure by using System Center to manage all SQL Servers. However, you cannot see any monitoring or management option for SQL Server in the System Center tools. What must you do?

**Answer:** Install the System Center Management Pack for SQL Server.



## Module Review and Takeaways

**Question:** What configuration requirements do you have for SQL server instances and databases in your organization?

**Answer:** Answers will vary, depending on students' experiences.

**Question:** Do you use the built-in best-practice policies in your workplace? If so, which ones have you imported?

**Answer:** Answers will vary, depending on students' experiences.



# Module 3

## Monitoring SQL Server 2014 Health

### Contents:

<b>Lesson 1:</b> Introduction to Health Monitoring	2
<b>Lesson 2:</b> Data Collector	4
<b>Lesson 3:</b> SQL Server Utility	7
Module Review and Takeaways	10
Lab Review Questions and Answers	11

## Lesson 1

# Introduction to Health Monitoring

### Contents:

Question and Answers	3
Resources	3

## Question and Answers

**Question:** You want to enable departmental database administrators (who are also responsible for the server and desktop operating systems) to view and assess simple performance data for the database instances they manage. Which health monitoring tool do you think is best suited to this task?

**Answer:** SQL Server Utility may be best suited because it provides simple diagnostic information and is relatively easy to configure.

## Resources

### Options for Monitoring Server Health in SQL Server 2014



**Reference Links:** You can download the System Center Monitoring Pack for SQL Server from the Microsoft download website.

## Lesson 2

# Data Collector

### Contents:

Question and Answers	5
Resources	5
Demonstration: Configuring Data Collector	5

## Question and Answers

**Question:** You want to enable the IT Junior Admins group access to reports generated from data collector information. Which of the following roles should you add the group to?

- ☐ mdw\_writer
- ☐ mdw\_reader
- ☐ dc\_admin
- ☐ dc\_operator
- ☐ dc\_proxy

**Answer:**

- ☐ mdw\_writer
- ☒ mdw\_reader
- ☐ dc\_admin
- ☐ dc\_operator
- ☐ dc\_proxy

## Resources

### Overview of Data Collector



**Reference Links:** For more information about the AMR tool, see *New AMR Tool: Simplifying the Migration to In-Memory OLTP* in the SQL Server Blog on the TechNet website.

## Demonstration: Configuring Data Collector

### Demonstration Steps

Configure the MDW

1. Ensure that the 20465C-MIA-DC, and 20465C-MIA-SQL virtual machines are running, and then log on to MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. In the D:\Demofiles\Mod03 folder, run **Setup.cmd** as Administrator.
3. Start SQL Server Management Studio and connect to the **MIA-SQL** database engine using Windows authentication.
4. In Object Explorer, expand **Management**, right-click **Data Collection**, click **Tasks**, and then click **Configure Management Data Warehouse**.
5. In the **Configure Management Data Warehouse Wizard**, click **Next**.
6. On the **Configure Management Data Warehouse Storage** page, in the **Server name** field, ensure that **MIA-SQL** is displayed, click **New**, in the **New Database** dialog box, in the **Database name** field, type **MDW**, click **OK**, and then click **Next**.
7. On the **Map Logins and Users** page, click **Next**, and then on the **Complete the Wizard** page, click **Finish**.
8. On the **Configure Data Collection Wizard Progress** page, click **Close**.
9. In Object Explorer, expand **Databases** and note that a database named **MDW** has been created.

#### Configure Data Collection

1. In Object Explorer, expand **Management**, right-click **Data Collection**, click **Tasks**, and then click **Configure Data Collection**.
2. In the **Configure Data Collection** wizard, click **Next**.
3. On the **Setup Data Collection Sets** page, to the right of the **Server name** field, click the ellipsis button, in the **Connect to Server** dialog box, type **MIA-SQL**, click **Connect**, and then in the **Database name** field, click **MDW**.
4. In the **Select data collector sets you want to enable** field, click the **System Data Collection Sets** check box, and then click **Next**.
5. On the **Complete the wizard** page, click **Finish**, and then when configuration is complete, click **Close**.

#### Upload Disk Usage Collection Sets

1. In Object Explorer, under **MIA-SQL**, under **Management**, expand **Data Collection**, expand **System Data Collection Sets**, and view the available collection sets.
2. Right-click **Disk Usage**, click **Collect and Upload Now**, and then when the **Collect and Upload Data Set** job completes, click **Close**.
3. Repeat the previous step for the **Query Statistics** and **Server Activity** collection sets.

#### View Management Data Warehouse Reports

1. In Object Explorer, right-click the **MDW** database, click **Reports**, point to **Management Data Warehouse**, and then click **Management Data Warehouse Overview**.
2. On the **Management Data Warehouse Overview: MDW** page, in the **MIA-SQL** row, click the link in the **Disk Usage** column.
3. On the **Disk Usage Collection Set** page, click **AdventureWorks**, and then on the **Disk Usage for database: AdventureWorks** page, review the disk usage statistics.
4. Close the **Disk Usage History** tab.
5. Close SQL Server Management Studio.



## Lesson 3

# SQL Server Utility

### Contents:

Question and Answers	8
Demonstration: Configuring a UCP	8

## Question and Answers

**Question:** You want to enroll a new server in your SQL Server Utility UCP. The UCP is in the IT.AdWorks.local domain. The new server is in the Sales.AdWorks.local domain. You receive an error when you try to complete the enrollment. How can you resolve this issue?

**Answer:** Check that there is a two-way trust relationship between the IT and Sales domains. This is required for the enrollment to work.

## Demonstration: Configuring a UCP

### Demonstration Steps

Create a UCP

1. Ensure that the 20465C-MIA-DC and 20465C-MIA-SQL virtual machines are running and then log on to MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. If you did not complete the previous demonstration in this module, in the D:\Demofiles\Mod03 folder, run **Setup.cmd** as Administrator.
3. Start SQL Server Management Studio and connect to the **MIA-SQL\SQL2** database engine instance using Windows authentication.
4. On the **View** menu, click **Utility Explorer**.
5. In Utility Explorer, on the **Getting Started** tab, click **Create a Utility Control Point (UCP)**.
6. In the **Create a Utility Control Point** wizard, on the **Introduction** page, review the information, and then click **Next**.
7. On the **Specify the Instance of SQL Server** page, click **Connect**.
8. In the **Connect to Server** dialog box, in the **Server name** box, type **MIA-SQL\SQL2**, and then click **Connect**.
9. On the **Specify the Instance of SQL Server** page, in the **Utility Control Point Name** box, note that the default name is **Utility**, and then click **Next**.
10. On the **Utility Collection Set Account** page, select **Use the SQL Server Agent service account** and click **Next**.
11. On the **SQL Server Instance Validation** page, wait for the validation operations to complete, and then click **Next**.
12. On the **Summary of UCP Creation** page, review the information, and then click **Next**.
13. On the **Utility Control Point Creation** page, wait for the creation operations to complete, and then click **Finish**.
14. In the **Utility Explorer Content** tab, note that there is **1** managed instance; but no data has been collected yet.
15. On the **Getting Started** tab, note the **Enroll instance of SQL Server with a UCP** link – you can use this to enroll additional servers in the UCP.

Run Jobs to Collect and Upload Data

1. In Object Explorer, under **MIA-SQL\SQL2**, expand **SQL Server Agent**, and then expand **Jobs**.
2. Right-click **sysutility\_mi\_collect\_performance**, and click **Properties**.
3. In the **Job Properties - sysutility\_mi\_collect\_performance** dialog box, on the **Schedules** page, note that this job is configured to run every 15 seconds. Then click **Cancel**.

**Note:** This job, and several others, is created when a server instance is enrolled in a UCP. The jobs run at scheduled times to collect and upload health metrics to the UCP server.

4. Right-click the **sysutility\_mi\_collect\_performance**, job and click **Start Job at Step**.
5. In the **Start Job – MIA-SQL\SQL2** dialog box, wait until the job completes, and then click **Close**.
6. Right-click **sysutility\_mi\_collect\_and\_upload**, and then click **Start Job at Step**.
7. In the **Start Job on 'MIA-SQL\SQL2'** dialog box, click **Start**, wait until the job completes, and then click **Close**.
8. Right-click **sysutility\_get\_views\_data\_into\_cache\_tables**, and then click **Start Job at Step**.
9. In the **Start Job on 'MIA-SQL\SQL2'** dialog box, click **Start**, wait until the job completes, and then click **Close**.
10. Right-click **sysutility\_get\_cache\_tables\_data\_into\_aggregate\_tables\_hourly**, and then click **Start Job at Step**.
11. In the **Start Job - MIA-SQL\SQL2** dialog box, wait until the job completes, and then click **Close**.
12. Right-click **sysutility\_get\_cache\_tables\_data\_into\_aggregate\_tables\_daily**, and then click **Start Job at Step**.
13. In the **Start Job on 'MIA-SQL\SQL2'** dialog box, click **Start**, wait until the job completes, and then click **Close**.

View the Dashboard

1. In the Utility Explorer pane, right-click **Utility (MIA-SQL\SQL2)** and click **Refresh**.
2. In the Utility Explorer pane, note that the **Managed Instance Health** chart now shows a single instance that is well utilized.

## Module Review and Takeaways

**Question:** With Data Collector, why is it better to have a central management data warehouse for data collection, than local installations?

**Answer:** To have a single point for performance analysis and to offload the storage from production systems.

**Question:** What challenges have you faced when planning to monitor server health and performing health monitoring? How have you attempted to overcome these problems?

**Answer:** Answers will vary, depending on students' experiences.

# Lab Review Questions and Answers

## Lab: Monitoring SQL Server Health

### Question and Answers

#### Lab Review

**Question:**

If the UCP reports CPU as consistently over-utilized, what steps would you take next to diagnose and resolve the issue?

**Answer:** Answers will vary depending on students' opinions, but the key thing to do is to monitor resource utilization by using tools such as Performance Monitor SQL Server Profiler to ascertain when over-utilization is occurring, and which activities are causing it.



# Module 4

## Consolidating Database Workloads with SQL Server 2014

### Contents:

<b>Lesson 1:</b> Considerations for Database Server Consolidation	2
<b>Lesson 2:</b> Managing Server Resources in a Consolidated Database Infrastructure	4
Module Review and Takeaways	10
Lab Review Questions and Answers	11

## Lesson 1

# Considerations for Database Server Consolidation

### Contents:

Question and Answers

3



## Question and Answers

**Question:** Your company has five different databases all hosted on different SQL Servers and installed by personnel in different departments. Recently, security breaches have been a problem caused by permissions mistakenly assigned by part-time administrators. The board have asked you to investigate database consolidation and to prioritize the centralized management of logins and permissions. What consolidation strategy should you recommend?

**Answer:** The requirement for simplified and centralized management of logins and permissions suggests using database-level consolidation.

## Lesson 2

# Managing Server Resources in a Consolidated Database Infrastructure

### Contents:

Question and Answers	5
Resources	5
Demonstration: Using Resource Governor	6
Demonstration: Configuring Instance-Level Resource Utilization	8

## Question and Answers

**Question:** You have six CPU cores in your SQL Server. You want to ensure that the two SQL Server instances never use the same core. You have executed the following T-SQL code on the first instance:

```
EXEC sp_configure 'show advanced options', 1;
```

```
RECONFIGURE;
```

```
GO
```

```
EXEC sp_configure 'affinity mask', 7;
```

```
RECONFIGURE;
```

```
GO
```

What affinity mask should you use on the second instance?

☐ 0

☐ 3

☐ 16

☐ 56

☐ 63

**Answer:**

☐ 0

☐ 3

☐ 16

☒ 56

☐ 63

## Resources

### Resource Management for an Instance of SQL Server by Using Resource Governor



**Reference Links:** For more information about Resource Governor and best practices for configuring it, download the white paper *Using the Resource Governor* from the MSDN library.

### Resource Management Across Multiple Instances of SQL Server



**Best Practice:** By default, the CPU affinity mask is not enabled, and a SQL Server instance can use all the available CPUs on a server as required. In most cases, this configuration is optimal, and you will not need to change affinity settings. To avoid negatively affecting performance, you should only configure CPU affinity and memory settings after very thorough consideration and testing.

## Resource Management in Virtualized Environments



**Reference Links:** For more information about System Center 2012 R2 Virtual Machine Manager, visit the Virtual Machine Manager page on the TechNet website.

### Demonstration: Using Resource Governor

#### Demonstration Steps

##### Create Resource Pools

1. Ensure that the 20465C-MIA-DC and 20465C-MIA-SQL virtual machines are both running, that you have logged on to 20465C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**, and that you have run **Setup.cmd** in the D:\Demofiles\Mod04 folder as Administrator.
2. Start SQL Server Management Studio, and connect to the **MIA-SQL** database engine using Windows authentication.
3. Open the **Resource Governor.sql** script from the D:\Demofiles\Mod04 folder.
4. In the query window, select the code under the comment **Create resource pools** and click **Execute**. This code creates two resource pools, named **Low Priority** and **High Priority**.

##### Create Workload Groups

1. In the query window, select the code under the comment **Create workload groups** and click **Execute**. This code creates a workload group named **DemoDB1WG** that uses the **Low Priority** resource pool, and a workload group named **DemoDB2WG** that uses the **High Priority** resource pool.
2. Select the code under the comment **Reconfigure resource governor** and click **Execute**. This reconfigures Resource Governor, enabling the resource pools and workload groups you have created.

##### Create a Classifier Function

1. In the query window, select the code under the comment **Create classifier function** and click **Execute**. This code creates a function that returns the appropriate workload group name for the current session, based on the name of the database to which the connection has been made.
2. Select the code under the comment **Add classifier function to resource governor** and click **Execute**. This reconfigures Resource Governor so that the function you created is used as the classifier function for all future connections.
3. Minimize SQL Server Management Studio.

##### Monitor Resource Pools and Workload Groups

1. Right-click the Start button and click **Computer Management**.
2. In Computer Management, in the pane on the left, expand **Performance**, expand **Monitoring Tools**, and then click **Performance Monitor**.
3. If any counters are listed under the chart, select them and press **Delete** so that the chart is blank.
4. On the toolbar, click **Add** (the green + icon).

5. In the **Add Counters** dialog box, in the list of objects, expand the **SQLServer: Resource Pool Stats** object, and then click **CPU control effect %**. Hold the CTRL key and click the following counters:
  - CPU usage target %
  - Max memory (KB)
  - Target memory (KB)
  - Used memory (KB)
6. In the **Instances of selected object** list click **High Priority**, hold the Ctrl key and click **Low Priority**, and then click **Add**. This adds the counters you selected for both resource pool instances.
7. In the **Add Counters** dialog box, in the list of objects, expand the **SQLServer:Workload Group Stats** object, and then click **CPU usage %**.
8. If the **Instances of selected object** list is empty, click **CPU usage %** again. In the **Instances of selected object** list, click **DemoDB1WG**, hold the Ctrl key and click **DemoDB2WG**, and then click **Add**.
9. In the **Add Counters** dialog box, click **OK**. Note that Performance Monitor displays the counter values. Click any of the counters under the chart and press **Ctrl+H**, and note that this highlights the currently selected counter in the graph.
10. Wait for the red line (which indicates the current time) to return to the beginning of the chart, and then in the D:\Demofiles\Mod04 folder, double-click **DemoDB1Query.cmd** to start a user query workload.
11. Observe the values of the counters in Performance Monitor until the red bar is approximately a third of the way across the chart.
12. With the DemoDB1Query.cmd command still running, in the D:\Demofiles\Mod04 folder, double-click **DemoDB2Query.cmd** to start the help desk workload. Observe the values of the counters in Performance Monitor until the red bar is approximately two thirds of the way across the chart.
13. Close the console window for the DemoDB2 query and observe the values of the counters in Performance Monitor until the red bar is almost all the way across the chart. Then in Performance Monitor, click the **Freeze Display** button before the red line reaches the end of the chart.
14. Close the console window for the DemoDB1Query query.
15. View the counters in the chart, and note the following:
  - The CPU control effect % counter shows the extent to which Resource Governor influenced CPU utilization.
  - The CPU target and actual usage for the DemoDB1 query was noticeably reduced during the period when the DemoDB2 query was running.
  - Neither resource group required its full allocation of memory – the workloads were CPU-intensive, but they were not memory-intensive.
16. Minimize Computer Management for the next demonstration.
17. In the D:\Demofiles\Mod04 folder, run **CleanUp.cmd**, as Administrator.

## Demonstration: Configuring Instance-Level Resource Utilization

### Demonstration Steps

#### Configure Server Settings

1. Ensure that you have completed the previous demonstration in this module.
2. In SQL Server Management Studio, in Object Explorer, on the **Connect** drop-down list, click **Database Engine**. When prompted, use Windows authentication to connect to the **MIA-SQL\SQL2** instance of the database engine.
3. In Object Explorer, right-click **MIA-SQL\SQL2** and click **Properties**.
4. In the **Server Properties – MIA-SQL\SQL2** dialog box, on the **Memory** page, change the **Maximum server memory (in MB)** value to 768.
5. On the **Processors** page, clear the **Automatically set processor affinity mask for all processors** check box, and then expand **All**, expand **NumaNode0**, and select the **Processor Affinity** check box for **CPU0**. Then click **OK**.
6. In Object Explorer, on the **Connect** drop-down list, click **Database Engine**. When prompted, use Windows authentication to connect to the **MIA-SQL\SQL3** instance of the database engine.
7. In Object Explorer, right-click **MIA-SQL\SQL3** and click **Properties**.
8. In the **Server Properties – MIA-SQL\SQL3** dialog box, on the **Memory** page, change the **Maximum server memory (in MB)** value to 512.
9. On the **Processors** page, clear the **Automatically set processor affinity mask for all processors** check box, and then expand **All**, expand **NumaNode0**, and select the **Processor Affinity** check box for **CPU1**. Then click **OK**.
10. Close SQL Server Management Studio without saving any files.

#### Monitor Instance Resource Utilization

1. Maximize Computer Management, and on the Performance Monitor node, under the chart, click the first counter, hold SHIFT and click the last counter, and press **Delete**.
2. On the toolbar, click **Add** (the green + icon).
3. In the **Add Counters** dialog box, in the list of objects, expand the **Processor** object, and then click **% Processor Time**. If the **Instances of selected object** list is empty, click **% Processor Time** again, and then click the **0** instance, hold the **CTRL** key and click the **1** instance, and then click **Add**.
4. In the **Add Counters** dialog box, in the list of objects, expand the **MSSQL\$SQL2: Memory Manager** object, and then click **Total Server Memory (KB)**. Then click **Add**.
5. In the **Add Counters** dialog box, in the list of objects, expand the **MSSQL\$SQL2: Resource Pool Stats** object, and then click **CPU usage %**. If the **Instances of selected object** list is empty, click **CPU usage %** again, and then select the **default** instance and click **Add**.
6. Repeat the previous two steps to add the same counters for the equivalent **MSSQL\$SQL3** objects.
7. In the **Add Counters** dialog box, click **OK**.
8. In Performance Monitor, click the **Unfreeze Display** button. Note that Performance Monitor displays the counter values.

9. In the D:\Demofiles\Mod04 folder, double-click **QueryMIA-SQL2.cmd**, and then view the counters in Performance Monitor. Note the following:
  - The **Total Server Memory (KB)** counter for the **MSSQL\$SQL2** instance rises to the 768 MB limit you set previously.
  - The **% Processor Time** counter for instance **0** rises in correlation with the **CPU usage %** counter for the **MSSQL\$SQL2** instance.
10. Close the command window for the MIA-SQL\SQL2 query.
11. In the D:\Demofiles\Mod04 folder, double-click **QueryMIA-SQL3.cmd**, and then view the counters in Performance Monitor. Note the following:
  - The **Total Server Memory (KB)** counter for the **MSSQL\$SQL3** instance rises to the 512 MB limit you set previously.
  - The **% Processor Time** counter for instance **1** rises in correlation with the **CPU usage %** counter for the **MSSQL\$SQL3** instance.
12. Close the command window for the MIA-SQL\SQL3 query and close Computer Management.
13. In the D:\Demofiles\Mod04 folder, run **CleanUp.cmd** as Administrator.

## Module Review and Takeaways

**Question:** Have you been involved in planning or implementing a consolidation initiative, or in managing SQL Server in a consolidated environment? If so, did the consolidation initiative succeed in delivering the intended benefits? How has consolidation affected administration? If you have not been involved in consolidation in any way, how do you think your organization might benefit from consolidating its SQL Server infrastructure?

**Answer:** Answers will vary, depending on individuals' experiences and opinions.



# Lab Review Questions and Answers

## Lab: SQL Server Consolidation

### Question and Answers

#### Lab Review

**Question:** Did you agree that the consolidation solution presented in the lab was the most appropriate one for the scenario? Would you have done anything differently, and if so, what?

**Answer:** Answers will vary, depending on individuals' opinions.



# Module 5

## Introduction to Cloud Data Solutions

### Contents:

Lesson 1: Overview of Cloud Computing	2
Lesson 2: SQL Server in a Private Cloud	4
Module Review and Takeaways	6

## Lesson 1

# Overview of Cloud Computing

### Contents:

Question and Answers

3

## Question and Answers

Categorize each item into the appropriate category. Indicate your answer by writing the category number to the right of each item.

Items	
1	Azure Database
2	Windows Azure
3	Office 365
4	Azure Pack
5	Hyper-V

Category 1		Category 2
Public Cloud Technologies		Private Cloud Technologies

**Answer:**

Category 1		Category 2
Public Cloud Technologies		Private Cloud Technologies
Windows Azure Azure Database Office 365		Azure Pack Hyper-V

## Lesson 2

# SQL Server in a Private Cloud

### Contents:

Question and Answers

5

## Question and Answers

**Question:** You have a SQL Server 2014 instance running on a physical server. You want to migrate this to a virtual server in your new private cloud. What tool should you use to complete the migration?

- ☐ The Windows sysprep utility
- ☐ The P2V Wizard
- ☐ SQL Server sysprep
- ☐ The MAP toolkit
- ☐ Windows Server failover clustering

**Answer:**

- ☐ The Windows sysprep utility
- ☒ The P2V Wizard
- ☐ SQL Server sysprep
- ☐ The MAP toolkit
- ☐ Windows Server failover clustering

## Module Review and Takeaways

**Question:** What benefits would a private cloud infrastructure bring to organizations that you have worked with?

**Answer:** Answers will vary, depending on students' experiences.



# Module 6

## Introduction to High Availability in SQL Server 2014

### Contents:

<b>Lesson 1:</b> High Availability Concepts and Options in SQL Server 2014	2
<b>Lesson 2:</b> Log Shipping	4
Module Review and Takeaways	10
Lab Review Questions and Answers	11

## Lesson 1

# High Availability Concepts and Options in SQL Server 2014

### Contents:

Question and Answers	3
Resources	3

## Question and Answers

**Question:** You have a website supported by a SQL Server that run Web Edition. As the website has recently taken on a more business-critical role, you want to increase the availability of the underlying database. Which of the following technologies should you use?

- ☐ Log Shipping
- ☐ Database Mirroring
- ☐ AlwaysOn Failover Cluster Instances
- ☐ AlwaysOn Availability Groups

**Answer:**

- ☒ Log Shipping
- ☐ Database Mirroring
- ☐ AlwaysOn Failover Cluster Instances
- ☐ AlwaysOn Availability Groups

## Resources

### Considerations for High Availability in a Private Cloud



**Reference Links:** For more information about guest clustering, including the storage options, see *Using Guest Clustering for High Availability* on the TechNet website.

## Lesson 2

# Log Shipping

### Contents:

Question and Answers	5
Resources	6
Demonstration: Configuring Log Shipping	6
Demonstration: Initiating Failover	8

## Question and Answers

Put the following planned failover steps in order by numbering each to indicate the correct order.

	Steps
	Back up the transaction log on the primary server by using the NORECOVERY option.
	Manually copy the log backup to the secondary server and then restore it, specifying the NORECOVERY option.
	Perform a database restore of the secondary database on the secondary server, specifying the RECOVERY option.
	Disable the backup, copy, and restore jobs on the primary and secondary servers.
	Copy any users and SQL Server Agent jobs to the secondary server from the primary server.
	Reconfigure applications that access the database to reach it by using the secondary server.
	Check the name of the latest backup on the primary server, and ensure that it has been applied to the database on the secondary server.
	Set up log shipping on the secondary server. Configure the old secondary server to be the new primary server and the old primary server to be the new secondary server.

**Answer:**

	Steps
2	Back up the transaction log on the primary server by using the NORECOVERY option.
3	Manually copy the log backup to the secondary server and then restore it, specifying the NORECOVERY option.
4	Perform a database restore of the secondary database on the secondary server, specifying the RECOVERY option.
5	Disable the backup, copy, and restore jobs on the primary and secondary servers.
7	Copy any users and SQL Server Agent jobs to the secondary server from the primary server.
8	Reconfigure applications that access the database to reach it by using the secondary server.
1	Check the name of the latest backup on the primary server, and ensure that it has been applied to the database on the secondary server.
6	Set up log shipping on the secondary server. Configure the old secondary server to be the new primary server and the old primary server to be the new secondary server.

## Resources

### SQL Server Agent and Log Shipping



**Reference Links:** For more information about service accounts for SQL Server Agent in log shipping scenarios, and about security on log shipping in general, see *How to Configure Security for SQL Server Log Shipping* on the Microsoft Support website.

### Considerations for Using Log Shipping



**Reference Links:** In addition to the alerts that you can use to monitor log shipping, there are various tables and stored procedures available to you. For more information about these tables and stored procedures, see *Monitor Log Shipping (Transact-SQL)* on the MSDN website.

### Pre-Installation Tasks




**Reference Links:** For more information about managing logins in log shipping, see *Management of Logins and Jobs After Role Switching* on the MSDN website.

## Demonstration: Configuring Log Shipping

### Demonstration Steps

#### Configure Log Shipping

1. Ensure that the 20465C-MIA-DC and 20465C-MIA-SQL virtual machines are running. Then log onto MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. In the D:\Demofiles\Mod06 folder, right-click **Setup.cmd** and click **Run as administrator**. When prompted, click **Yes**.
3. Start SQL Server Management Studio and connect to MIA-SQL by using Windows authentication.
4. In Object Explorer, expand **Databases**, right-click **DemoDB**, click **Tasks**, and then click **Ship Transaction Logs**.
5. In the **Database Properties** window, on the **Transaction Log Shipping** page, click the **Enable this as a primary database in a log shipping configuration** check box.
6. Click **Backup Settings**. In the **Network path to backup folder** field, type **\\MIA-SQL\PrimaryBackupFolder**.
7. In the **If the backup folder is located on the primary server, type a local path to the folder** field, type **D:\DemoFiles\Mod06\PrimaryBackupFolder**.
8. Click **Schedule**, in the **New Job Schedule** dialog box, under **Daily frequency**, in the **Occurs every** field, type **3**, and then click **OK**. This configures the job to run every three minutes.
9. In the **Transaction Log Backup Settings** window, click **OK**.
10. In the **Database Properties** window, click **Add**.
11. In the **Secondary Database Settings** window, click **Connect**, in the **Connect to Server** dialog box, in the **Server name** field, type **MIA-SQL\SQL2**, and then click **Connect**.

12. On the **Initialize Secondary Database** tab, ensure that **Yes, generate a full backup of the primary database and restore it into the secondary database** is selected.
13. Click the **Copy Files** tab. In the **Destination folder for copied files** field, type **\\MIA-SQL\SecondaryRestoreFolder**.
14.  **Note:** In this demonstration, we are using two instances of SQL Server on the same Windows server. In a real world scenario, you would normally use two separate Windows servers.
15. Click **Schedule**, in the **New Job Schedule** dialog box, under **Daily frequency**, in the **Occurs every** field, type **3**, and then click **OK**.
16. Click the **Restore Transaction Log** tab, and then click **Schedule**. In the **New Job Schedule** dialog box, under **Daily frequency**, in the **Occurs every** field, type **3**, and then click **OK**.
17. Click **Standby mode**, and then click the **Disconnect users in the database when restoring backups** check box.
18. In the **Secondary Database Settings** window, click **OK**, and then in the **Database Properties – DemoDB** window, click **OK**.
19. In the **Save Log Shipping Configuration** dialog box, wait for the configuration to complete, and then click **Close**.

Inspect the Log Shipping Configuration

1. In Object Explorer, expand **SQL Server Agent**, expand **Jobs**, and note the two log shipping jobs, which are called **LSAlert\_MIA-SQL** and **LSBackup\_DemoDB**.
2. In Object Explorer, click **Connect**, click **Database Engine**, in the **Connect To Server** dialog box, in the **Server name** field, type **MIA-SQL\SQL2**, and then click **Connect**.
3. In Object Explorer, under **MIA-SQL\SQL2**, expand **SQL Server Agent**, expand **Jobs**, and note the three log shipping jobs, which are called **LSAlert\_MIA-SQL\SQL2**, **LSCopy\_MIA-SQL\_DemoDB**, and **LSRestore\_MIA-SQL\_DemoDB**.
4. Expand **Databases**, and note that the **DemoDB** database is now present on the MIA-SQL\SQL2 instance, and that it shows as **Standby / Read-Only**.
5. In Object Explorer, click **MIA-SQL**.
6. On the taskbar, click **Windows Explorer**, browse to **D:\Demofiles\Mod06\PrimaryBackupFolder**, and note the **DemoDB.Bak** full backup and the one or more transaction log backups that are in the folder. If there are no transaction log backups present, wait for up to three minutes and then check again.
7. Browse to **D:\Demofiles\Mod06\SecondaryRestoreFolder**, and note the transaction log backup or backups that were copied from the primary server. If there are no transaction log backups present, wait for up to three minutes and then check again.

Test Access to Log Shipping Servers

1. In SQL Server Management Studio, in Object Explorer, click **MIA-SQL**.
2. In Windows Explorer, browse to **D:\Demofiles\Mod06**, and double-click **TestPrimary.sql**.
3. In the **TestPrimary.sql** query window, under the comment **Query a table in the DemoDB database**, highlight the Transact-SQL statements and then click **Execute**.
4. Review the results, and note that the value in the **Color** column is **Black**.

5. Under the comment **Update the color value to Red for ProductKey 210**, highlight the Transact-SQL statements and then click **Execute**.
  6. Review the results, and note that the value in the **Color** column is now **Red**.
  7. In Object Explorer, under MIA-SQL, under SQL Server Agent, right-click **LSBackup\_DemoDB**, and then click **Start Job at Step**.
  8. In the **Start Jobs - MIA-SQL** dialog box, wait for the steps to complete, and then click **Close**.
  9. In Object Explorer, under MIA-SQL\SQL2, under SQL Server Agent, right-click **LSCopy\_MIA-SQL\_DemoDB**, and then click **Start Job at Step**.
  10. In the **Start Jobs - MIA-SQL\SQL2** dialog box, wait for the steps to complete, and then click **Close**.
  11. Repeat steps 9 and 10 to run the **LSRestore\_MIA-SQL\_DemoDB** job.
  12. In Windows Explorer, in the **D:\Demofiles\Mod06** folder, double-click **TestSecondary.sql**.
  13. In the **TestSecondary.sql** query window, under the comment **Query a table in the DemoDB database**, highlight the Transact-SQL statements and then click **Execute**.
  14. Review the results, and note that the value in the **Color** column is **Red**, indicating that the change you made on the primary server has been restored on the secondary server.
  15. Under the comment **Update the color value to Blue for ProductKey 210**, highlight the Transact-SQL statement and then click **Execute**.
  16. Review the message stating that it is not possible to update the read-only database.
  17. Close the **TestPrimary.sql** and **TestSecondary.sql** query windows, and do not save any changes.
- Leave SQL Server Management Studio open for the next demonstration.

## Demonstration: Initiating Failover

### Demonstration Steps

Initiate Manual Failover

1. Ensure that you have completed the previous demonstration in this module.
2. In Object Explorer, click **MIA-SQL**, and then open **BackupLog.sql** in the **D:\Demofiles\Mod06** folder.
3. In SQL Server Management Studio, in the **BackupLog.sql** query window, below the comment **Backup the DemoDB log on the primary server**, highlight the Transact-SQL statement, and then click **Execute**.
4. In Windows Explorer, in the **D:\Demofiles\Mod06\PrimaryBackupFolder** folder, right-click **FinalBackup.trn**, and then click **Copy**. Browse to **D:\Demofiles\Mod06\SecondaryRestoreFolder**, right-click in the folder, and then click **Paste**.
5. In SQL Server Management Studio, in Object Explorer, click **MIA-SQL\SQL2**. Then open **RestoreLog.sql** in the **D:\Demofiles\Mod06** folder.
6. In Object Explorer, under **MIA-SQL\SQL2**, under SQL Server Agent, right-click **LSCopy\_MIA-SQL\_DemoDB**, and then click **Start Job at Step**.
7. In the **Start Jobs - MIA-SQL\SQL2** dialog box, wait for the steps to complete, and then click **Close**.
8. Repeat steps 6 and 7 to run the **LSRestore\_MIA-SQL\_DemoDB** job.
9. In SQL Server Management Studio, in the **RestoreLog.sql** query window, under the comment **Restore the final log backup**, highlight the Transact-SQL statement, and then click **Execute**.



10. Under the comment **Recover the database**, highlight the Transact-SQL statement, and then click **Execute**.
11. In Object Explorer, under MIA-SQL, right-click **DemoDB**, and then click **Refresh**. Note that the DemoDB database is now in a **Restoring** state.
12. In Object Explorer, under **MIA-SQL\SQL2**, right-click **DemoDB**, and then click **Refresh**. Note that the **DemoDB** database is now online.
13. Open **TestFailover.sql** in the D:\Demofiles\Mod06 folder.
14. In the **TestFailover.sql** query window, under the comment **View a record**, highlight the Transact-SQL statement, and then click **Execute**.
15. In the **TestFailover.sql** query window, under the comment **Update a record**, highlight the Transact-SQL statement, and then click **Execute**.
16. Close SQL Server Management Studio, and do not save any changes.
17. In Windows Explorer, in the D:\Demofiles\Mod06 folder, right-click **Cleanup.cmd**, click **Run as Administrator**, and when prompted, click **Yes**.

## Module Review and Takeaways

**Question:** In what scenarios might an organization choose log shipping as a high availability solution?

**Answer:** Answers will vary, depending on students' experiences.

# Lab Review Questions and Answers

## Lab: Using Log Shipping

### Question and Answers

#### Lab Review

**Question:** If you wanted to fail back to the original primary server, what would you need to do?

**Answer:** Disable the existing log shipping jobs, and then set up log shipping so that the original primary server was configured as a secondary server, and the original secondary server was configured as the primary server.



# Module 7

## Clustering with Windows Server and SQL Server 2014

### Contents:

<b>Lesson 1:</b> Introduction to Windows Server Failover Clustering	2
<b>Lesson 2:</b> SQL Server AlwaysOn Failover Cluster Instances	5
Module Review and Takeaways	10
Lab Review Questions and Answers	11

## Lesson 1

# Introduction to Windows Server Failover Clustering

### Contents:

Question and Answers	3
Resources	3

## Question and Answers

Put the following steps in an updating run in order by numbering each to indicate the correct order.


	Steps
	The update coordinator installs the required updates and restarts the node if necessary.
	The update coordinator puts the first node into maintenance mode.
	The update coordinator moves any clustered roles that the node hosts to another node in the cluster.
	The update coordinator moves the clustered roles back to the node.
	The update coordinator removes the node from maintenance mode before moving on to the next node.

**Answer:**


	Steps
3	The update coordinator installs the required updates and restarts the node if necessary.
1	The update coordinator puts the first node into maintenance mode.
2	The update coordinator moves any clustered roles that the node hosts to another node in the cluster.
4	The update coordinator moves the clustered roles back to the node.
5	The update coordinator removes the node from maintenance mode before moving on to the next node.


## Resources

### Overview of Windows Server Failover Clustering

 **Reference Links:** Windows Server 2012 R2 includes many enhancements to WSFC. For a list of these enhancements, see the *What's New in Failover Clustering in Windows Server 2012 R2?* page on the TechNet website.

### Storage for Windows Server Failover Clusters

 **Reference Links:** For more information about using storage spaces with WSFC, see *How to Configure a Clustered Storage Space in Windows Server 2012* on the MSDN blogs website.

 **Reference Links:** For a list of the enhancements to SMB in SMB 3.02, see the *What's New in SMB in Windows Server 2012 R2?* page on the TechNet website.

## Cluster-Aware Updating



**Reference Links:** For more information about Cluster-Aware Updating, see *Cluster-Aware Updating Overview* on the TechNet website. For more information about using Cluster-Aware Updating to apply software updates to SQL Server FCIs, download the technical article *Patching SQL Server Failover Cluster Instances with Cluster-Aware Updating* from the TechNet website.



## Lesson 2

# SQL Server AlwaysOn Failover Cluster Instances

### Contents:

Question and Answers	6
Resources	6
Demonstration: Implementing an AlwaysOn Failover Cluster Instance	6
Demonstration: Configuring Failover Options	9

## Question and Answers

**Question:** You want to ensure that when the **sp\_server\_diagnostics** stored procedure returns the value **resource error**, failover is initiated in AlwaysOn FCI cluster. What failover condition level should you choose?

- ( ) Level 1
- ( ) Level 2
- ( ) Level 3
- ( ) Level 4
- ( ) Level 5

**Answer:**

- ( ) Level 1
- ( ) Level 2
- ( ) Level 3
- (√) Level 4
- ( ) Level 5

## Resources

### Enhancements to AlwaysOn Failover Cluster Instances



**Reference Links:** For more information about the **sp\_server\_diagnostics** system stored procedure, see the *sp\_server\_diagnostics (Transact SQL)* topic in SQL Server Books Online. For more information about failover policies, see the *Failover Policy for Failover Cluster Instances* topic in SQL Server Books Online.

## Demonstration: Implementing an AlwaysOn Failover Cluster Instance

### Demonstration Steps

Prepare for Cluster Installation

1. Ensure that only the 20465C-MIA-DC, 20465C-MIA-FCI-CLUST1, 20465C-MIA-FCI-CLUST2, and 20465C-MIA-FCI-CLUST3 virtual machines are running, and log on to 20465C-MIA-FCI-CLUST1 as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. In the D:\Demofiles\Mod07 folder, right-click **Setup.cmd**, and then click **Run as administrator**.
3. In the **User Account Control** dialog box, click **Yes**, and wait for the script to finish.
4. In the D:\Demofiles\Mod07 folder, right-click **Setup2.cmd**, and then click **Run as administrator**.
5. In the **User Account Control** dialog box, click **Yes**, and wait for the script to finish.
6. On the Start screen, type **Failover** and then start the **Failover Cluster Manager** app.

7. Click **MIA-CLUSTER.adventureworks.msft**, and then review the information about the Windows Server Failover Cluster, noting the following points:
  - The cluster includes three nodes.
  - The cluster has no clustered roles.
  - The quorum configuration is Node Majority.
8. Expand **MIA-CLUSTER.adventureworks.msft**, and then click **Nodes**.
9. In the **Nodes** pane, in the **Status** column, note that the status of each of the three cluster nodes is **Up**.
10. Expand **Storage**, and then click **Disks**. In the **Disks** pane, in the **Status** column, note that the status of **Cluster Disk 1** is **Online**.
11. Right-click **Cluster Disk 1**, and then click **Properties**.
12. In the **Cluster Disk 1 Properties** dialog box, on the **General** tab, review the information about the cluster disk.
13. Click the **Policies** tab, and review the failure responses for the disk resource.
14. Click the **Advanced Policies** tab, and review the possible owners for the disk resource.
15. In the **Cluster Disk 1 Properties** dialog box, click **Cancel**.
16. Click **Networks**, and then in the **Networks** pane, in the **Status** column, note that the status of the **Cluster Network 1** resource is **Up**.
17. Click **MIA-CLUSTER.adventureworks.msft**, and then click **Validate Cluster**.
18. In the **Validate a Configuration Wizard**, on the **Before you Begin** page, click **Next**.
19. On the **Testing Options** page, click **Run all tests (recommended)**, and then click **Next**.
20. On the **Review Storage Status** page, select the **Cluster Disk 1** check box, and then click **Next**.
21. On the **Confirmation** page, click **Next**.
22. Wait for the validation to complete, and then click **View Report**.
23. In the **Failover Cluster Validation Report**, in the **Cluster Configuration** section, click **Validate Quorum Configuration**, review the information, and then click **Back to Failover Cluster Validation Report**.
24. In the **Network** section, click **Validate Network Communication**, review the information, and then click **Back to Failover Cluster Validation Report**.
25. Close the **Cluster Failover Validation Report**, and then in the **Validate a Configuration Wizard**, click **Finish**.
26. Close Failover Cluster Manager.

#### Install a New Failover Cluster Instance

1. In the C:\SQLServer2014-x64-ENU folder, double-click **Setup.exe**, and then in the **User Account Control** dialog box, click **Yes**.
2. Wait a few moments for the **SQL Server Installation Center** to start, click **Installation**, and then click **New SQL Server failover cluster installation**.
3. Wait a few moments for the **Install a SQL Server Failover Cluster** wizard to start. On the **Global Rules** page, wait for the rule check to complete. Then on the **Microsoft Updates** and **Product Updates** pages, clear any checkboxes and click **Next**.

4. After the **Install Setup Files** page completes, on the **Install Failover Cluster Rules** page, wait for the rules check to complete, review the results, and then click **Next**.
5. On the **Product Key** page, click **Next**, on the **License Terms** page click the **I accept the license terms** check box, clear the option to turn on Customer Experience Improvement Program ("CEIP") and Error Reporting, and then click **Next**.
6. On the **Setup Role** page, click the **SQL Server Feature Installation** radio button, and then click **Next**.
7. On the **Feature Selection** page, click the **Database Engine Services** check box, and then click **Next**.
8. On the **Instance Configuration** page, in the **SQL Server Network Name** field, type **DEMOSQLCLUSTER**, click **Named instance**, in the **Named instance** field, type **DEMOSQL**, and then click **Next**.
9. On the **Cluster Resource Group** page, review the information, and then click **Next**.
10. On the **Cluster Disk Selection** page, ensure that the **Cluster Disk 1** check box is selected, and then click **Next**.
11. On the **Cluster Network Configuration** page, click the **IP Type** check box, in the **Address** column, type **10.10.0.160**, and then click **Next**.
12. On the **Server Configuration** page, in the **SQL Server Agent** row, in the **Account Name** column, type **ADVENTUREWORKS\ServiceAcct**, in the **Password** column, type **Pa\$\$w0rd**, in the **SQL Server Database Engine** row, in the **Account Name** column, type **ADVENTUREWORKS\ServiceAcct**, in the **Password** column, type **Pa\$\$w0rd**, and then click **Next**.
13. On the **Database Engine Configuration** page, on the **Server Configuration** tab, click **Add Current User**, and then click **Next**.

Add a Node to the Failover Cluster Instance

1. Log on to the **MIA-CLUST2** virtual machine as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. In the C:\SQLServer2014-x64-ENU folder, double-click **Setup.exe**, and then in the **User Account Control** dialog box, click **Yes**.
3. In SQL Server Installation Center, click **Installation**, and then click **Add node to a SQL Server failover cluster**.
4. On the **Global Rules** page, wait for the rule check to complete. Then on the **Microsoft Updates** and **Product Updates** pages, clear any checkboxes and click **Next**.
5. After the **Install Setup Files** page completes, on the **Add Node Rules** page, wait for the rules check to complete, review the results, and then click **Next**.
6. On the **Product Key** page, click **Next**, on the **License Terms** page click the **I accept the license terms** check box, clear the **Turn on Customer Experience Improvement Program and Error Reporting** check box, and then click **Next**.
7. On the **Cluster Node Configuration** page, note that the node will join the **DEMOSQL** Failover Cluster Instance, click **Next**, and then on the **Cluster Network Configuration** page, click **Next**.
8. On the **Service Accounts** page, in the **SQL Server Agent** row, in the **Password** column, type **Pa\$\$w0rd**, in the **SQL Server Database Engine** row, in the **Password** column, type **Pa\$\$w0rd**, and then click **Next**.
9. On the **Ready to Add Node** page, click **Install**. Wait for the installation to complete, on the **Complete** page, click **Close**, close the SQL Server Installation Center, and then close File Explorer.

### Review the Failover Cluster Configuration

1. On MIA-CLUST1, on the Start screen, type **Failover**, and then click **Failover Cluster Manager**.
2. In Failover Cluster Manager, click **MIA-CLUSTER.adventureworks.msft**, and then click **Roles**.
3. In the **Roles** pane, click **SQL Server (DEMOSQL)**, review the information in the **Status** and **Owner Node** columns, in the **Actions** pane, click **Properties**, review the properties of the role, and then click **Cancel**.
4. In Failover Cluster Manager, at the bottom of the **SQL Server (DEMOSQL)** pane, click the **Resources** tab, and then review the information about the resources associated with SQL Server (DEMOSQL).
5. Minimize **Failover Cluster Manager**. Then start SQL Server Management Studio and connect to the **DEMOSQLCLUSTER\DEMOSQL** database engine instance using Windows authentication.
6. In Object Explorer, right-click **DEMOSQLCLUSTER\DEMOSQL**, click **Properties**, and then in the **Server Properties – DEMOSQLCLUSTER\DEMOSQL** dialog box, on the **General** tab, in the **IsClustered** row, note that the value is **True**.
7. Click **Cancel**, and then minimize SQL Server Management Studio.

## Demonstration: Configuring Failover Options

### Demonstration Steps

#### Trigger Failover

1. Ensure that you have completed the previous demonstration in this module.
2. On the MIA-CLUST2 virtual machine, on the Start screen type **Failover** and then start **Failover Cluster Manager**.
3. In Failover Cluster Manager, expand **MIA-CLUSTER.adventureworks.msft**, click **Roles**, and then in the **Owner Node** column, note that the current owner of the role is **MIA-CLUST1**.
4. On the host computer, in **Hyper-V Manager**, right-click **20465C-MIA-FCI-CLUST1**, and then click **Shut Down**. In the **Shut Down Machine** dialog box, click **Shut Down**. This action simulates failure of the node that is the owner of the SQL Server (DEMOSQL) role.
5. On the MIA-CLUST2 virtual machine, in **Failover Cluster Manager**, click **Roles**, and then watch the SQL Server (DEMOSQL) role failover from the MIA-CLUST1 node to the MIA-CLUST2 node.

When failover has finished, in the lower part of the window, click the **Resources** tab, and review the available resources, checking that they all report a status of **Online**.

## Module Review and Takeaways

### Review Question(s)

**Question:** Can you think of any high availability scenarios in your organization that you might implement from new or upgrade by using an AlwaysOn FCI?

**Answer:** Answers will vary, depending on the students' opinions and experiences.

# Lab Review Questions and Answers

## Lab: Implementing an AlwaysOn Failover Cluster Instance

### Question and Answers

#### Lab Review

**Question:** In the lab, you configured the FCI to automatically fail back to a preferred node immediately the node became available. Why might you decide not to do this in some situations?

**Answer:** You might not configure an FCI to automatically fail back to a preferred node immediately so that you can perform checks and tests on the node before allowing it to take ownership of any cluster resources again.





# Module 8

## AlwaysOn Availability Groups

### Contents:

<b>Lesson 1:</b> Introduction to AlwaysOn Availability Groups	2
<b>Lesson 2:</b> Working with AlwaysOn Availability Groups	4
<b>Lesson 3:</b> Considerations for Using AlwaysOn Availability Groups	12
Module Review and Takeaways	15
Lab Review Questions and Answers	16

## Lesson 1

# Introduction to AlwaysOn Availability Groups

### Contents:

Question and Answers	3
Resources	3

## Question and Answers

**Question:** You have configured a flexible failover policy for your AlwaysOn availability group and you have set the health check timeout threshold to 24 seconds. How often will the WSFC call the **sp\_server\_diagnostics** stored procedure?

- ☐ Once every 2.4 seconds
- ☐ Once every 12 seconds
- ☐ Once every 24 seconds
- ☐ Once every 8 seconds
- ☐ Once every 72 seconds

**Answer:**

- ☐ Once every 2.4 seconds
- ☐ Once every 12 seconds
- ☐ Once every 24 seconds
- ☒ Once every 8 seconds
- ☐ Once every 72 seconds

## Resources

### Active Secondary Replicas



**Reference Links:** You will learn more about read-only routing in the next lesson.

### Failover in an AlwaysOn Availability Group



**Reference Links:** For more information about failover in AlwaysOn Availability Groups, see the *Failover and Failover Modes (AlwaysOn Availability Groups)* topic in SQL Server Books Online.

### Flexible Failover Policies



**Reference Links:** For more information about the lease that the availability group uses to connect to the WSFC, see the *How It Works: SQL Server AlwaysOn Lease Time-out* article on the MSDN website.



**Reference Links:** For more information about configuring flexible failover policies for AlwaysOn Availability Groups, see the *Configure the Flexible Failover Policy to Control Conditions for Automatic Failover (AlwaysOn Availability Groups)* article on the MSDN website.

## Lesson 2

# Working with AlwaysOn Availability Groups

### Contents:

Question and Answers	5
Resources	5
Demonstration: Creating and Testing an AlwaysOn Availability Group	6
Demonstration: Implementing and Testing Read-Only Routing	10

## Question and Answers

Put the following data synchronization steps in order by numbering each to indicate the correct order.

	Steps
	On the primary replica, cached log data is flushed to the log file on disk and the log records are prepared for queuing.
	The secondary replicas receive the messages from the primary replica and place them into cache.
	The log pages on the secondary replica are applied to the secondary database data file.
	On the primary replica, the messages in the queues are sent to the secondary replicas.
	The log is applied to the secondary database and an acknowledgement sent to the primary replica.
	The log for each database is captured and sent to the corresponding queue ready to be copied to the secondary replicas.

**Answer:**

	Steps
1	On the primary replica, cached log data is flushed to the log file on disk and the log records are prepared for queuing.
4	The secondary replicas receive the messages from the primary replica and place them into cache.
6	The log pages on the secondary replica are applied to the secondary database data file.
3	On the primary replica, the messages in the queues are sent to the secondary replicas.
5	The log is applied to the secondary database and an acknowledgement sent to the primary replica.
2	The log for each database is captured and sent to the corresponding queue ready to be copied to the secondary replicas.

## Resources

### Enhancements to AlwaysOn Availability Groups in SQL Server 2014



**Reference Links:** For more information about using AlwaysOn Availability Groups to provide high availability for In-Memory OLTP databases, see the article *In-Memory OLTP: High Availability for Databases with Memory-Optimized Tables* on the SQL Server Blog, on the TechNet website.



**Reference Links:** For more information about the **sys.fn\_hadr\_is\_primary\_replica** system function, see the article *sys.fn\_hadr\_is\_primary\_replica (Transact-SQL)* in SQL Server Books Online.

## Pre-Requisites for Using AlwaysOn Availability Groups



**Reference Links:** For more information about hotfixes to support AlwaysOn Availability Groups on WSFCs, see the article *Prerequisites, Restrictions, and Recommendations for AlwaysOn Availability Groups (SQL Server)* on the MSDN website.

## Pre-Installation Tasks



**Reference Links:** For more information about Transact-SQL statements and PowerShell cmdlets you can use to configure an Availability Group, see SQL Server Books Online.

## Read-Only Routing



**Reference Links:** For more information about configuring read-only routing by using PowerShell, see the article *Configure Read-Only Routing for an Availability Group (SQL Server)* on the TechNet website.

## Managing AlwaysOn Availability Groups



**Reference Links:** For more information about monitoring AlwaysOn Availability Groups by using PowerShell cmdlets, see the article *Monitoring AlwaysOn Health with PowerShell* on the SQL AlwaysOn Team Blog website.

## Monitoring AlwaysOn Availability Groups



**Reference Links:** For more information about the DMVs and system functions that you can use to monitor AlwaysOn Availability Groups, see the *AlwaysOn Availability Groups Dynamic Management Views and Functions (Transact-SQL)* webpage in SQL Server Books Online.



**Reference Links:** For more information about the errors that trigger automatic page repair, see the article *Automatic Page Repair (Availability Groups/Database Mirroring)* on the SQL Server Books Online website.

## Demonstration: Creating and Testing an AlwaysOn Availability Group

### Demonstration Steps

Create an AlwaysOn Availability Group

1. Ensure that only the 20465C-MIA-DC, 20465C-MIA-AG-CLUST1, 20465C-MIA-AG-CLUST2, and 20465C-MIA-AG-CLUST3 virtual machines are running, and then log on to 20465C-MIA-AG-CLUST1 as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. In the D:\Demofiles\Mod08 folder, right-click **Setup.cmd**, and then click **Run as administrator**.
3. In the **User Account Control** dialog box, click **Yes**, wait for the script to finish.

4. Start SQL Server Management Studio and connect to **MIA-CLUST1** using Windows authentication.
5. In Object Explorer, expand **Databases**, right-click the **DemoDB1**, point to **Tasks**, and then click **Back Up**.
6. In the **Back Up Database – DemoDB1** dialog box, in the **Destination** list, click the existing backup file path, click **Remove**, and then click **Add**.
7. In the Select Backup Destination dialog box, in the File name box, type **D:\Demofiles\Mod08\DemoShare\DemoDB1.bak**, and then click **OK**.
8. In the **Back Up Database – DemoDB1** dialog box, in the **Backup type** list, ensure that **Full** is selected, and then click **OK**.
9. In the Microsoft SQL Server Management Studio dialog box, click **OK**.
10. Repeat steps 5 to 9 to perform a full database backup of the **DemoDB2** database, using the backup path **D:\Demofiles\Mod08\DemoShare\DemoDB2.bak**.
11. In SQL Server Management Studio, in Object Explorer, expand **AlwaysOn High Availability**, right-click **Availability Groups**, and then click **New Availability Group Wizard**.
12. In the New Availability Group wizard, on the Introduction page, click **Next**.
13. On the Specify Availability Group Name page, in the Availability group name box, type **Demo-AG**, and then click **Next**.
14. On the **Select Databases** page, select the **DemoDB1** and **DemoDB2** database check boxes, and then click **Next**.
15. On the Specify Replicas page, on the Replicas tab, click **Add Replica**.
16. In the **Connect to Server** dialog box, in the **Server name** box, type **MIA-CLUST2**, in the **Authentication** list, ensure **Windows Authentication** is selected, and then click **Connect**.
17. Repeat steps 15 and 16 to add **MIA-CLUST3** as a replica.
18. On the **Replicas** tab, select the **Automatic Failover** check box for **MIA-CLUST1** and **MIA-CLUST2**. This automatically selects the **Synchronous Commit** check box for these replicas.
19. On the **Replicas** tab, in the **Readable Secondary** list for **MIA-CLUST2**, click **Read-intent only**.
20. On the **Replicas** tab, in the **Readable Secondary** list for **MIA-CLUST3**, click **Yes**.
21. Review the default settings on the **Endpoints** and **Backup Preferences** tabs, and then click the **Listener** tab.
22. Click **Create an availability group listener**, and then specify the following settings:
  - Listener DNS Name: **MIA-DEMO-CLUST**
  - **Port**: **1433**
  - Network Mode: **Static IP**
23. Click **Add**, and in the **Add IP Address** dialog box, in the **IPv4 Address** box, type **10.10.0.55**, and then click **OK**.
24. **Note:** If **Add** is not visible, you may need to maximize the dialog box or increase the screen resolution of the virtual machine).
25. On the **Specify Replicas** page, click **Next**.
26. On the **Select Initial Data Synchronization** page, ensure that **Full** is selected.

27. In the Specify a shared network location accessible by all replicas box, type \\MIA-CLUST1\DemoShare, and then click Next.
28. On the **Validation** page, review the validation results, and then click **Next**.
29. On the **Summary** page, click **Finish**.

Test Connectivity to the Availability Group

1. Right-click the Start button, and then click **Command Prompt**.
2. At the command prompt, type the following command to open a SQLCMD session and connect to the MIA-DEMO-CLUST availability group listener, and then press Enter:

```
sqlcmd -E -S MIA-DEMO-CLUST
```

3. At the command prompt, type the following commands to verify that the SQLCMD session is connected to the primary replica (MIA-CLUST1):

```
SELECT @@ServerName  
GO
```

4. At the command prompt, type the following commands to retrieve rows from the **Employee** table in the **DemoDB1** database, and then view the results:

```
SELECT * FROM DemoDB1.dbo.Employee  
GO
```

5. At the command prompt, type the following command to exit the SQLCMD session, and then press Enter:

```
Exit
```

6. At the command prompt, type the following command to open a SQLCMD session and connect to the MIA-CLUST3 replica, and then press Enter:

```
sqlcmd -E -S MIA-CLUST3
```

7. At the command prompt, type the following commands to retrieve rows from the **Orders** table in the **DemoDB2** database, and then view the results:

```
SELECT * FROM DemoDB2.dbo.Orders  
GO
```

8. At the command prompt, type the following command to exit the SQLCMD session, and then press Enter:

```
Exit
```

9. At the command prompt, type the following command to open a SQLCMD session and connect to the MIA-CLUST2 replica, and then press Enter:

```
sqlcmd -E -S MIA-CLUST2
```



10. At the command prompt, type the following commands to attempt to retrieve rows from the **Employee** table in the **DemoDB1** database:

```
SELECT * FROM DemoDB1.dbo.Employee
GO
```

11. View the error that is returned.
12. At the command prompt, type the following command to exit the SQLCMD session, and then press Enter:

```
Exit
```

13. At the command prompt, type the following command to open a SQLCMD session and connect to the MIA-CLUST2 replica with a read-intent connection, and then press Enter:

```
sqlcmd -E -S MIA-CLUST2 -K ReadOnly
```

14. At the command prompt, type the following commands to attempt to retrieve rows from the **Employee** table in the **DemoDB1** database, and then view the results:

```
SELECT * FROM DemoDB1.dbo.Employee
GO
```

15. At the command prompt, type the following command to exit the SQLCMD session, and then press Enter:

```
Exit
```

#### Change the Primary Replica

1. In SQL Server Management Studio, in Object Explorer, under **AlwaysOn High Availability** and **Availability Groups**, right-click **Demo-AG (Primary)**, and then click **Show Dashboard**.
2. In the upper-right corner of the dashboard, click **Start Failover Wizard**.
3. In the Fail Over Availability Group: Demo-AG wizard, on the Introduction page, click **Next**.
4. On the **Select New Primary Replica** page, note the warning about data loss if you decide to choose MIA-CLUST3 as the new primary replica. Select the **MIA-CLUST2** check box, and then click **Next**.
5. On the **Connect to Replica** page, click **Connect**.
6. In the **Connect to Server** dialog box, click **Connect**.
7. On the **Connect to Replica** page, click **Next**.
8. On the **Summary** page, click **Finish**.
9. On the **Results** page, click **Close**.
10. In the dashboard, note that the primary instance is now MIA-CLUST2 and that MIA-CLUST1 is a secondary replica. Note that the dashboard may take a few seconds to update after failover, so if the new configuration doesn't display immediately, just wait until the dashboard updates.
11. On the Taskbar, maximize the Command Prompt window.
12. At the command prompt, type the following command to open a SQLCMD session and connect to the MIA-DEMO-CLUST availability group listener, and then press Enter:

```
sqlcmd -E -S MIA-DEMO-CLUST
```

13. At the command prompt, type the following commands to verify that the SQLCMD session is connected to the new primary replica (MIA-CLUST2):

```
SELECT @@ServerName
GO
```

14. At the command prompt, type the following command to exit the SQLCMD session, and then press Enter:

```
Exit
```

15. Close the command prompt.
16. Repeat steps 1 to 9 to move the primary replica back to MIA-CLUST1.
17. Keep SQL Server Management Studio open for the next demonstration.

## Demonstration: Implementing and Testing Read-Only Routing

### Demonstration Steps

Configure read-only routing

1. Ensure that you have completed the previous demonstration in this module.
2. In SQL Server Management Studio, open the **ReadOnlyRouting.sql** script file in the D:\Demofiles\Mod08 folder.
3. In the query window, under the comment **Alter the Availability Group to add read-only routing URLs**, highlight the Transact-SQL statements, and then click **Execute**.
4. Under the comment **Configure read-only routing lists, which control how each replica behaves when it is the primary replica**, highlight the Transact-SQL statements, and then click **Execute**.

Test read-only routing

1. In the lower left corner of the task bar, click **Start**, type **cmd**, and then click **Command Prompt**.
2. At the command prompt, type the following to connect to the listener, and then press Enter:

```
sqlcmd -E -S MIA-DEMO-CLUST -d DemoDB1 -K ReadOnly
```

3. At the command prompt, type the following, and then press Enter:

```
SELECT @@Servername Enter
GO
```

Note that the result returned is **MIA-CLUST2** (the primary replica is MIA-CLUST1).

4. At the command prompt, type the following, and then press Enter:

```
Exit
```

5. At the command prompt, type the following to connect to the listener, and then press Enter:

```
sqlcmd -E -S MIA-DEMO-CLUST -d DemoDB1
```

6. At the command prompt, type the following, and then press Enter:

```
SELECT @@Servername Enter
GO
```

7. Note that the result returned is **MIA-CLUST1** because the `-K` parameter was not used in the connection string.
8. At the command prompt, type the following, and then press Enter:

```
Exit
```

9. Close the command prompt.

## Lesson 3

# Considerations for Using AlwaysOn Availability Groups

### Contents:

Question and Answers	13
Resources	13
Demonstration: Viewing Quorum Configuration for the Availability Group	
Host Cluster	13

## Question and Answers

**Question:** You have been asked to estimate the failover time for an AlwaysOn availability cluster. You have obtained the following values:

Health check timeout threshold: 30 seconds

Redo queue size on secondary servers (averaged over 1 hour at peak times): 10 transactions

Typical redo rate on secondary servers: 5 transactions per second

Failover overhead time: 45 seconds

What do you estimate the failover time to be for this cluster?

**Answer:** 1 min and 17 seconds

## Resources

### Scenarios for Using Availability Groups



**Best Practice:** After failover, you should ensure that you configure a high availability solution at the disaster recovery site to ensure that access to the services can continue uninterrupted.

### Failover Time and Data Loss



**Reference Links:** For more information about monitoring and troubleshooting AlwaysOn Availability Groups, see the *Monitor Performance for AlwaysOn Availability Groups* webpage on the MSDN website.

## Demonstration: Viewing Quorum Configuration for the Availability Group Host Cluster

### Demonstration Steps

View Quorum Configuration

1. In SQL Server Management Studio, in the Demo-AG Availability Group dashboard, click **View Cluster Quorum Information**.
2. In the **Cluster Quorum Information** dialog box, review the information, and note that each of the three nodes in the cluster has one vote.
3. In the **Cluster Quorum Information** dialog box, click **Close**.

Use the Configure Quorum Cluster Wizard

1. On the Start screen, type **Failover** and then start **Failover Cluster Manager**.
2. Click **MIA-CLUSTER.adventureworks.msft**, in the **Actions** pane, click **More Actions**, and then click **Configure Cluster Quorum Settings**.
3. In the **Configure Cluster Quorum Wizard** dialog box, click **Next**.
4. On the **Select Quorum Configuration Option** page, click **Advanced quorum configuration**, and then click **Next**.

5. On the **Select Voting Configuration** page, click **Select Nodes**, clear the **MIA-CLUST3** check box, and then click **Next**.
6. On the **Select Quorum Witness** page, click the **Configure a file share witness** radio button, and then click **Next**.
7. On the **Configure File Share Witness** page, in the **File Share Path** field, type **\\MIA-DC\DemoWitnessShare**, and then click **Next**.
8. On the **Confirmation** page, click **Cancel** to cancel the change in the configuration.
9. Close SQL Server Management Studio, and then close Failover Cluster Manager.
10. On the taskbar, click **File Explorer**, browse to **D:\Demofiles\Mod08**, right-click **Cleanup.cmd**, click **Run as administrator**, in the **User Account Control** dialog box, click **Yes**, and then close File Explorer.

## Module Review and Takeaways

**Question:** In what scenarios might an organization use AlwaysOn Availability Groups to implement high availability?

**Answer:** Answers will vary, depending on students' opinions.

## Lab Review Questions and Answers

### Lab: Implementing an Always On Availability Group

#### Question and Answers

##### Lab Review

**Question:** What was the purpose of removing the quorum vote from the MIA-CLUST3 cluster node?

**Answer:** To prevent the node from influencing the cluster quorum configuration. This is necessary because the failure of this node should not affect the other cluster nodes.



# Module 9

## Planning High Availability and Disaster Recovery

### Contents:

<b>Lesson 1:</b> High Availability and Disaster Recovery with SQL Server 2014	2
<b>Lesson 2:</b> SQL Server High Availability and Disaster Recovery Solutions	4
Module Review and Takeaways	6
Lab Review Questions and Answers	7

## Lesson 1

# High Availability and Disaster Recovery with SQL Server 2014

### Contents:

Question and Answers	3
Resources	3

## Question and Answers

**Question:** You are configuring an AlwaysOn availability group to provide both high availability and disaster recovery. The availability group includes 4 servers at your main datacenter and 2 servers in the disaster recovery site at a different physical location. To which of the following servers should you **not** give a quorum vote?

- ☐ A server in the disaster recovery site that hosts a secondary replica.
- ☐ A server at the main datacenter that hosts a primary replica
- ☐ A server at the main datacenter that hosts a secondary replica
- ☐ A fileshare witness

**Answer:**

- ☒ A server in the disaster recovery site that hosts a secondary replica.
- ☐ A server at the main datacenter that hosts a primary replica
- ☐ A server at the main datacenter that hosts a secondary replica
- ☐ A fileshare witness

## Resources

### Considerations for Implementing High Availability and Disaster Recovery



**Reference Links:** For more information about forcing quorum, see the article *Force a WSFC Cluster to Start Without a Quorum* on the MSDN website.



**Reference Links:** For general guidelines on the key features and requirements of the various SQL Server 2014 HA and DR technologies and the differences between them, review the topic *Planning High Availability*, in the module *Introduction to High Availability in SQL Server 2014*, in this course.

## Lesson 2

# SQL Server High Availability and Disaster Recovery Solutions

### Contents:

Question and Answers

5

## Question and Answers

**Question:** You have a SQL Server that uses the Simple recovery model. You want to ensure availability. Should you create an AlwaysOn FCI or an AlwaysOn Availability Group?

**Answer:** Create an AlwaysOn FCI

## Module Review and Takeaways

**Question:** Think about how high availability and disaster recovery implementations work in your own organization. Can you think of any ways in which you could improve the solutions that are currently in place?

**Answer:** Answers will vary, depending on students' opinions and experiences.

# Lab Review Questions and Answers

## Lab: Planning High Availability and Disaster Recovery

### Question and Answers

#### Lab Review

**Question:** In the lab exercise, you considered various solutions for each scenario. Do you agree with the recommended solution for each scenario? If not, what would you have done differently?

**Answer:** Answers will vary, depending on students' opinions.





# Module 10

## Replicating Data

### Contents:

Lesson 1: SQL Server Replication	2
Lesson 2: Planning Replication	7
Module Review and Takeaways	9
Lab Review Questions and Answers	10

## Lesson 1

# SQL Server Replication

### Contents:

Question and Answers	3
Resources	3
Demonstration: Implementing Replication	4

## Question and Answers

**Question:** You are configuring a replication between two servers and you want to add the second server as a subscriber to a publication on the publishing server. Which of the following roles must you be a member of?

- ☐ db\_owner
- ☐ public
- ☐ replmonitor
- ☐ sysadmin
- ☐ db\_datareader

**Answer:**

- ☐ db\_owner
- ☐ public
- ☐ replmonitor
- ☒ sysadmin
- ☐ db\_datareader

## Resources

### Data Validation and Monitoring



**Reference Links:** For more information about performing data validation, see the article *Validate Data at the Subscriber* in the MSDN library.

### Troubleshooting Replication



**Reference Links:** You can also use the **tablediff** command prompt utility to compare replicated tables and identify mismatching data. For more information about the **tablediff** utility, see the article *How to: Compare Replicated Tables for Differences (Replication Programming)* in the MSDN library.



**Reference Links:** For more information about verifying the status of replication agents, see the articles *How to: View Information and Perform Tasks for the Agents Associated with a Publication (Replication Monitor)* and *How to: View Information and Perform Tasks for the Agents Associated with a Subscription (Replication Monitor)* in the MSDN library.



**Reference Links:** For more information about memory problems due to replication agents, see the article *Running a Large Number of Agents is Causing Memory Problems* in the MSDN library.



**Reference Links:** For more information about publishing stored procedure replication, see the article *Publishing Stored Procedure Execution in Transactional Replication* in the MSDN library.



**Reference Links:** For more information about performance planning for replication, see the article *Enhancing General Replication Performance* in the MSDN library.



**Reference Links:** For more information about security issues and replication, see the article *Security Issues are Preventing Data from Being Replicated* in the MSDN library.



**Reference Links:** For more information about troubleshooting replication in general, see the *Troubleshooting Concepts (Replication)* article in the MSDN library.

## Demonstration: Implementing Replication

### Demonstration Steps

Create a Publication

1. Ensure that the 20465C-MIA-DC and 20465C-MIA-SQL virtual machines are running and then log on to 20465C-MIA-SQL as **ADVENTUREWORKS\Student** with the password **Pa\$\$w0rd**.
2. In the D:\Demofiles\Mod10 folder, run **Setup.cmd** as Administrator.
3. Start SQL Server Management Studio and connect to the **MIA-SQL** database engine instance using Windows authentication.
4. In SQL Server Management Studio, in Object Explorer, expand **Replication**, right-click **Local Publications**, and then click **New Publication**.
5. In the **New Publication Wizard**, on the **New Publication Wizard** page, click **Next**.
6. On the **Distributor** page, ensure that the '**MIA-SQL**' will act as its own **Distributor** option is selected, and then click **Next**.
7. On the **Snapshot Folder** page, review the default location for the snapshot folder, and then click **Next**.
8. On the **Publication Database** page, click **DemoDB**, and then click **Next**.
9. On the **Publication Type** page, click **Transactional publication**, and then click **Next**.
10. On the **Articles** page, in the **Objects to publish** box, expand **Tables**, select the **Product (Products)** check box, and then click **Next**.
11. On the **Filter Table Rows** page, click **Next**.
12. On the **Snapshot Agent** page, select the **Create a snapshot immediately and keep the snapshot available to initialize subscriptions** check box, and then click **Next**.
13. On the **Agent Security** page, click **Security Settings**.
14. In the **Snapshot Agent Security** dialog box, in the **Process account** box, type **ADVENTUREWORKS\ServiceAcct**, in the **Password** and **Confirm Password** boxes, type **Pa\$\$w0rd**, click **OK**, and then on the **Agent Security** page, click **Next**.
15. On the **Wizard Actions** page, ensure that only **Create the publication** is selected, and then click **Next**.
16. On the **Complete the Wizard** page, in the **Publication name** box, type **DemoPub**, and then click **Finish**.
17. On the **Creating Publication** page, wait for the operation to complete, and then click **Close**.

### Create a Subscription

1. In Object Explorer, click **Connect Object Explorer**.
2. In the **Connect to Server** dialog box, in the **Server name** box, type **MIA-SQL\SQL2**, and then click **Connect**.
3. In Object Explorer, under **MIA-SQL\SQL2**, expand **Replication**, right-click **Local Subscriptions**, and then click **New Subscriptions**.
4. In the **New Subscription Wizard**, on the **New Subscription Wizard** page, click **Next**.
5. On the **Publication** page, in the **Publisher** list, click **Find SQL Server Publisher**.
6. In the **Connect to Server** dialog box, in the **Server name** box, type **MIA-SQL**, and then click **Connect**.
7. On the **Publication** page, click **DemoPub**, and then click **Next**.
8. On the **Distribution Agent Location** page, ensure that **Run each agent at its Subscriber (pull subscriptions)** is selected, and then click **Next**.
9. On the **Subscribers** page, in the **Subscription Database** column, click **New database**.
10. In the **New Database** dialog box, in the **Database name** box, type **DemoDB**, and then click **OK**.
11. On the **Subscribers** page, click **Next**.
12. On the **Distribution Agent Security** page, click the ellipsis button.
13. In the **Distribution Agent Security** dialog box, in the **Process account** box, type **ADVENTUREWORKS\ServiceAcct**, in the **Password** and **Confirm Password** boxes, type **Pa\$\$w0rd**, and then click **OK**.
14. On the **Distribution Agent Security** page, click **Next**.
15. On the **Synchronization Schedule** page, ensure that **Agent Schedule** is set to **Run continuously**, and then click **Next**.
16. On the **Initialize Subscriptions** page, ensure that **Immediately** is selected, and then click **Next**.
17. On the **Wizard Actions** page, ensure that only **Create the subscription(s)** is selected, and then click **Next**.
18. On the **Complete the Wizard** page, review the configuration steps, and then click **Finish**.
19. On the **Creating Subscription(s)** page, wait for the operation to complete, and then click **Close**.

### Test Replication

1. In Object Explorer, under **MIA-SQL\SQL2**, expand **Databases**, expand **DemoDB**, expand **Tables**, right-click **Products.Product**, and then click **Select Top 1000 Rows**.
2. In the **Results** pane, review the information, noting that the query returned data for **Product 1**, **Product 2**, and **Product 3**.
3. In Object Explorer, under **MIA-SQL**, expand **Databases**, right-click **DemoDB**, and click **New Query**.
4. In the query window, type the following Transact-SQL statement, and then click **Execute**:

```
INSERT INTO Products.Product
VALUES
('Product 4');
GO
```

5. Wait for 30 seconds, and then switch to the query tab you used to select the top 1000 rows from the **Products.Product**, in the **MIA-SQL\SQL2** instance, and click Execute.
6. In the **Results** pane, review the information, noting that the table now includes data for **Product 4**.
7. Close SQL Server Management Studio, and do not save any changes.
8. In the D:\Demofiles\Mod10 folder, run **Cleanup.cmd** as Administrator.

## Lesson 2

# Planning Replication

### Contents:

Question and Answers

8

## Question and Answers

**Question:** You have created a publication that consists of the staff list from the HumanResources database. You want to ensure that HR personnel can quickly locate staff information wherever they are in your global company. What type of replication should you use?

**Answer:** Answers may vary, but since this is a list that only occasionally changes, depending on your staff turnover, you should probably consider using snapshot replication.



## Module Review and Takeaways

**Question:** What are the main reasons you can think of for implementing SQL Server replication?

**Answer:** To implement load balancing, improve availability, reduce network traffic, or enable offline processing.

# Lab Review Questions and Answers

## Lab: Planning and Implementing Replication

### Question and Answers

#### Lab Review

**Question:** In the lab, the replication of changes at the subscriber to the publisher was triggered manually. What could you configure so that the changes are automatically replicated? How might you configure this feature for the **HumanResources** database?

**Answer:** You could configure a schedule to automatically replicate changes from the subscriber to the publisher. In the **HumanResources** database, users update their timesheet data after each shift, so you could schedule replication to occur a short time afterwards.