DBMoto®

Microsoft SQL Server Always On Configuration
Version 9.0.0.10

Software Release Date: 3/15/16

Document History

Version	Author	Date	Reviewer	Date	Approver	Date	Comments
1	JHLorenzin	6/7/16	VFarruggio	6/7/16	VFarruggio	6/7/16	

Table of Contents

Overview

This document gives detailed steps on setting up transactional replications on a Microsoft SQL Server database that is part of an Always On Availability Group. It uses the environment described below as an example.

For a full explanation of Availability Groups and the Always On concept, refer to Microsoft SQL Server documentation, for example [Overview of Always On Availability Groups (SQL Server)](#). Here is a brief overview as it relates to setting up Always On Availability Groups with DBMoto.

In a Microsoft SQL Server environment, an *availability group* can provide a failover environment for a set of user databases by supporting a set of primary databases and one to eight sets of secondary databases. Each set of availability database is hosted by an *availability replica*. A single *primary replica* hosts the primary databases, and one to eight *secondary replicas*, which host a set of secondary databases that serve as a potential failover targets for the availability group. The use of Always On Availability Groups requires the establishment of a Windows Server Failover Clustering (WSFC) cluster. Each availability replica within an availability group resides on a different node of the same WSFC cluster.

Environment

The example used to explain how to set up and use an Always On Availability Group with DBMoto is described below.

The initial Always On SQL Server environment:

VH4-206: Synchronous Replica   – Current Primary
VH4-207: Synchronous Replica
Availability Group: TEST1_AG
AG database: AGTest1
AG Listener: TEST1_AG_Listen   (192.168.1.81, port 1433)
Cluster name: AGTestcluster
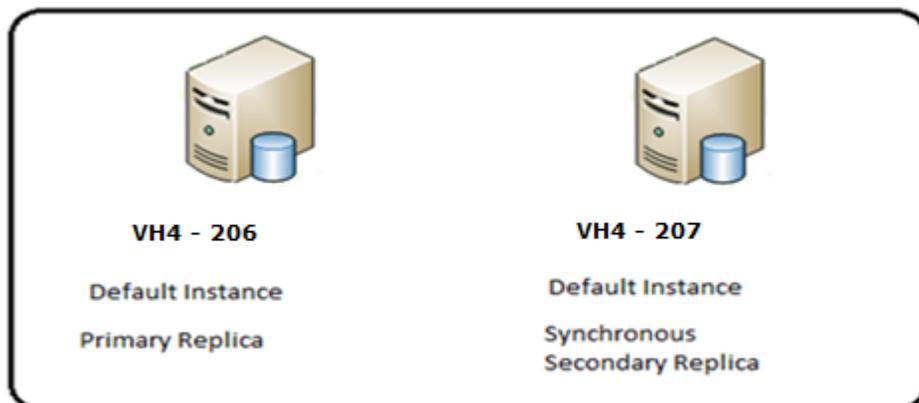
The environment set up for replication with DBMoto:

VH4-206: Original Publisher
VH4-207: Publisher Replica
VH4-205: Distributor

NOTE: Do not set a distributor on any of the publishers in this case as the failover of a distributor is not supported.
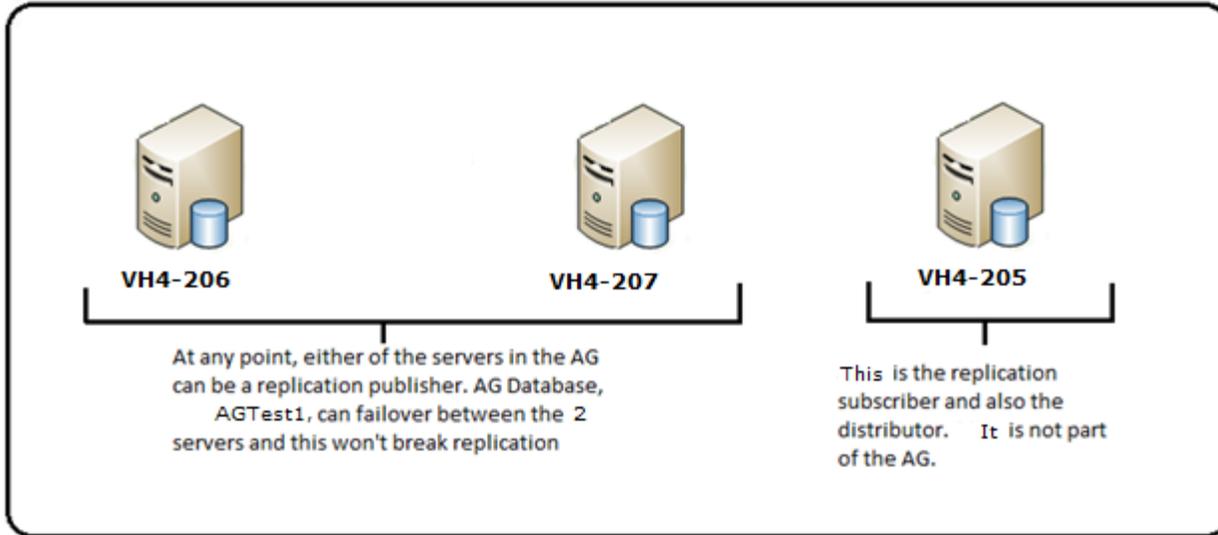
Configure the Replication Environment

1. Configure a Remote Distributor

To connect to VH4-205 from MS SQL Server Management Studio:

1. Right click on **Replication** and select **Configure Distribution**.
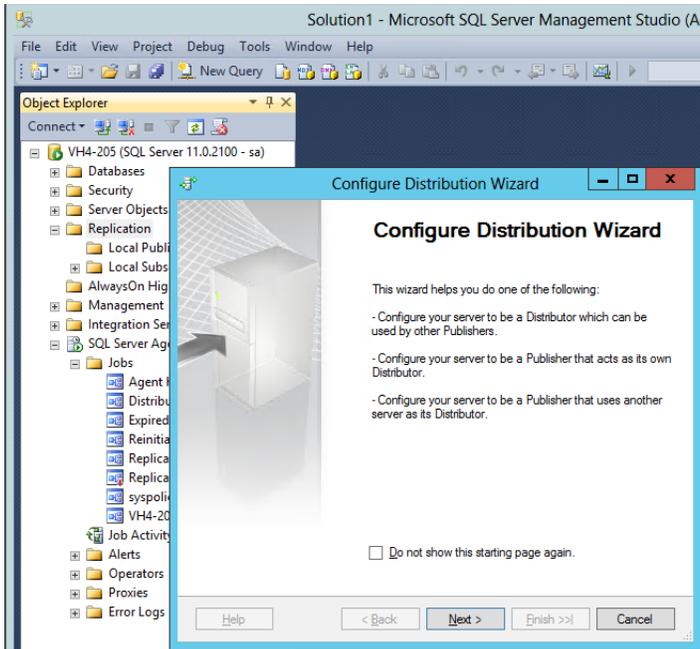
2. Select the first option to set up VH4-205 as distributor:
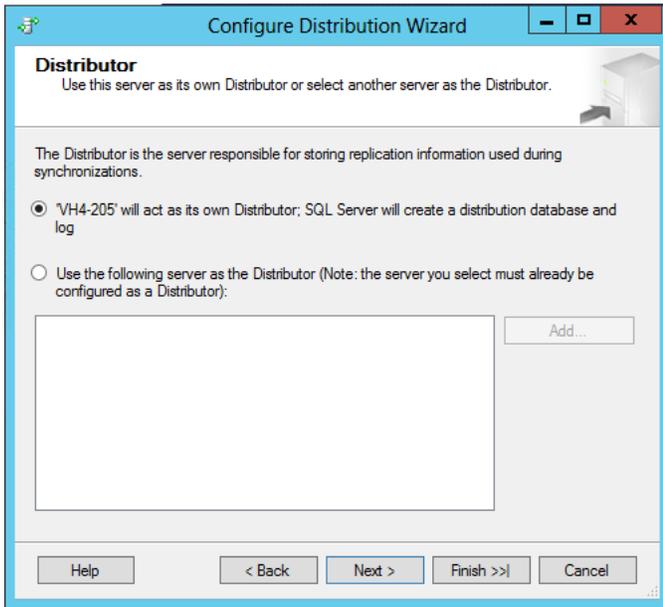
3. Click **Next**.

4. The Server Agent should be started automatically.
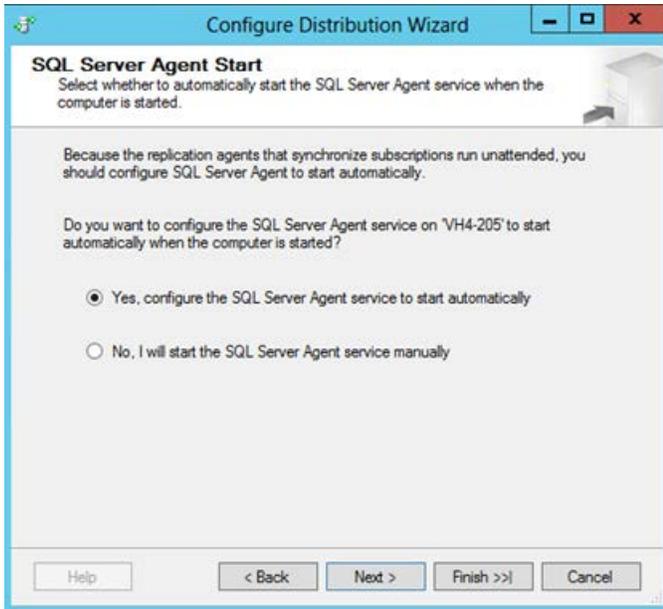
5. Click **Next**.

6. Specify the **Snapshot Folder** location:
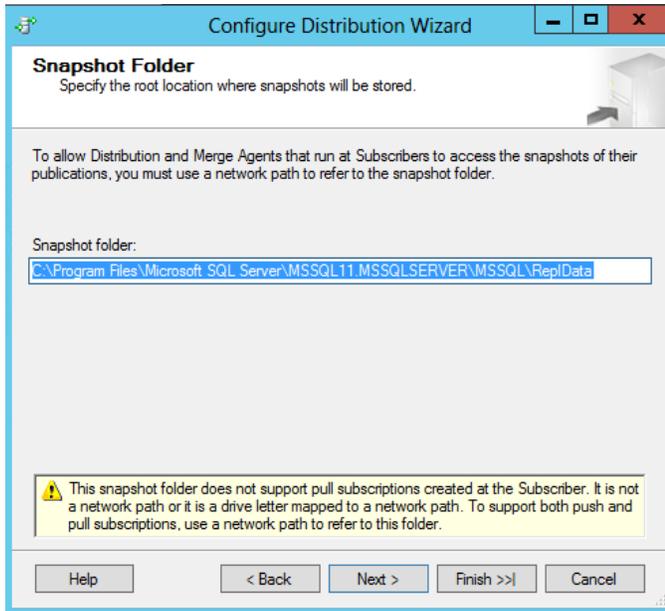
7. Click **Next**.

8. Name the distributor 'DBRS_distribution' and specify the path for the database file and database log file:
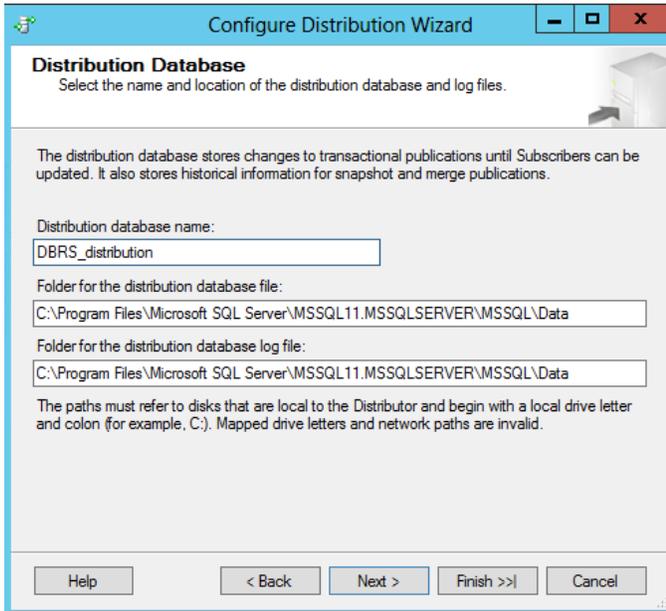
9. Click **Next** to specify VH4-205, VH4-206 and VH4-207 as publishers.
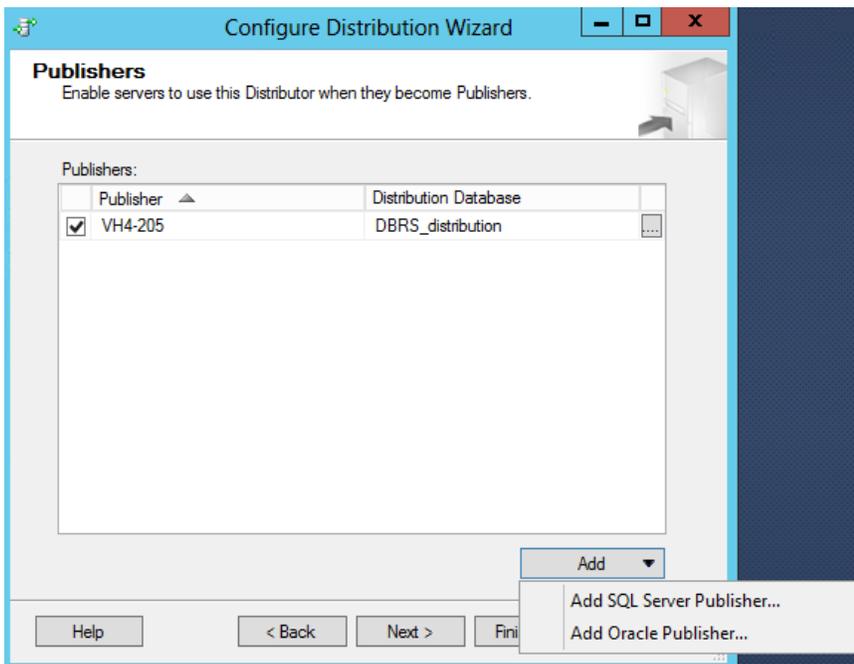
10. Click **Add**, then **Add SQL Server Publisher** to connect to each server that will act as publisher. Note that VH4-205 already exists in the list and you can choose to leave it that way.

11. Check that your screen looks like the one below after adding the servers.
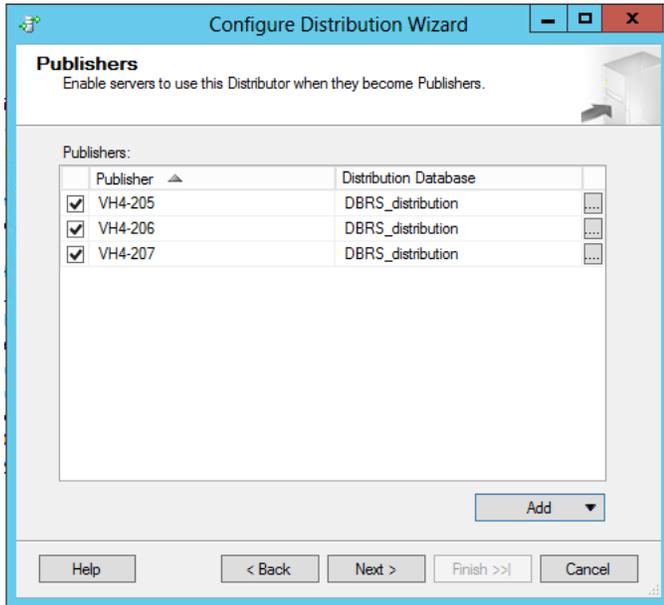
12. Click **Next**.

13. Enter the password that the publishers will use to connect to the distributor:
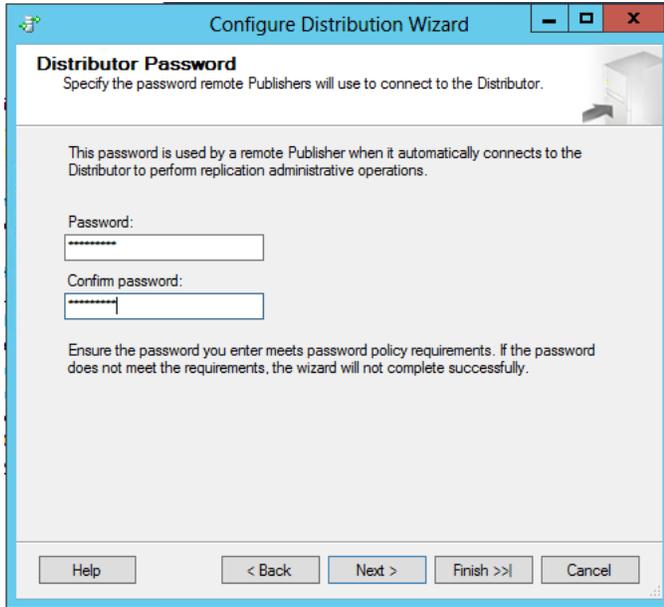
> **NOTE:** This password is used internally between publishers and distributor. Make a note of it as you will need to use it again when configuring a new publisher.

14. Click **Next**, then select **Configure distribution**.

15. Click **Next**.

16. Click **Finish** to complete the wizard.

Now the distributor is successfully set up.

2. Configure Distribution at the Primary Replica

1. Connect to the primary replica, VH4-206:

2. Right click on **Replication** and select **Configure Distribution**:
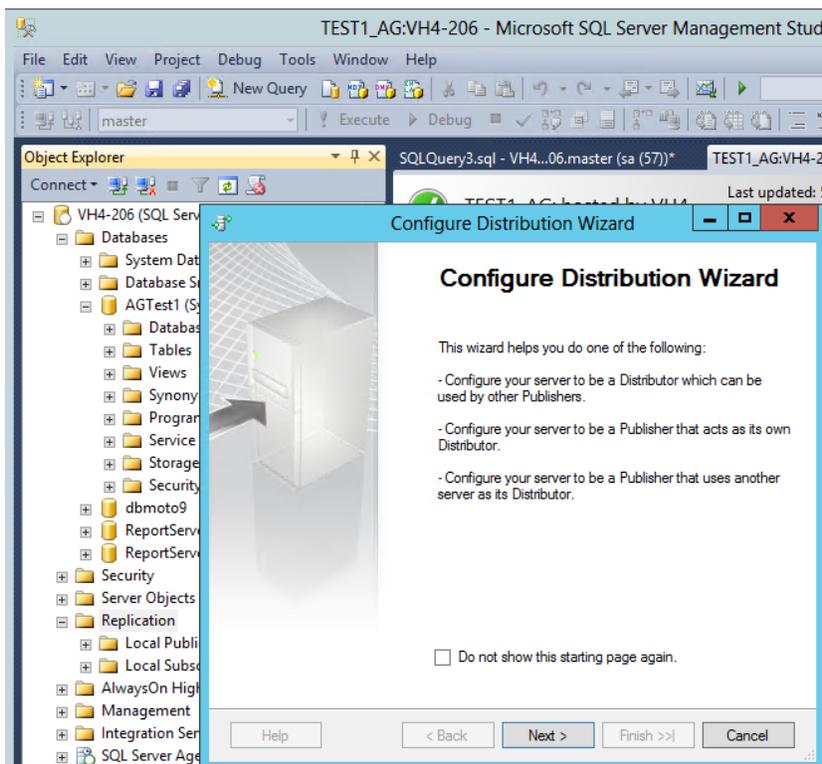
3. Select the option to connect to another server as the remote distributor.

4. Click **Add** and select VH4-205.
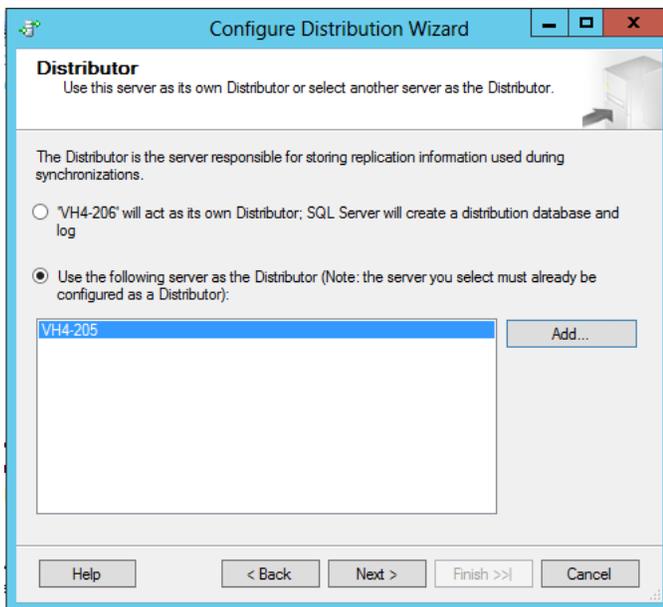
5. Type the same password used earlier to configure the distributor:
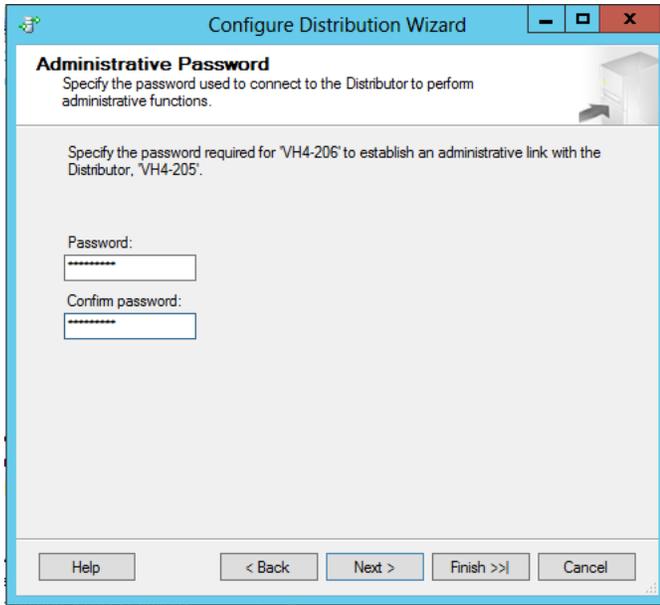
6. Click **Next**.

7. Select **Configure distribution**.

8. Click **Next**.

9. Click **Finish** to complete the wizard.

Now the remote distributor is successfully set up.

3. Configure Distribution for the Secondary Replica

To connect to the secondary replica, VH4-207, repeat the same steps as for the primary replica.

4. Set Log Reader Agent PublisherFailoverPartner Property

The SQL Server Log Reader Agent is set up to monitor the transaction log on the primary publisher, VH4-206. In case of failover to VH4-207, the agent job cannot be switched to pick up changes from the secondary (now primary) replica, unless the PublisherFailoverPartner property is set to VH4-207. This is the failover partner instance participating in a database mirroring session with the publication database. On failover, the publisher of the secondary replica will start to replicate to the remote distributor.

1. Connect to the distributor, VH4-205.

2. Select **Replication** and select **Distributor Properties**.
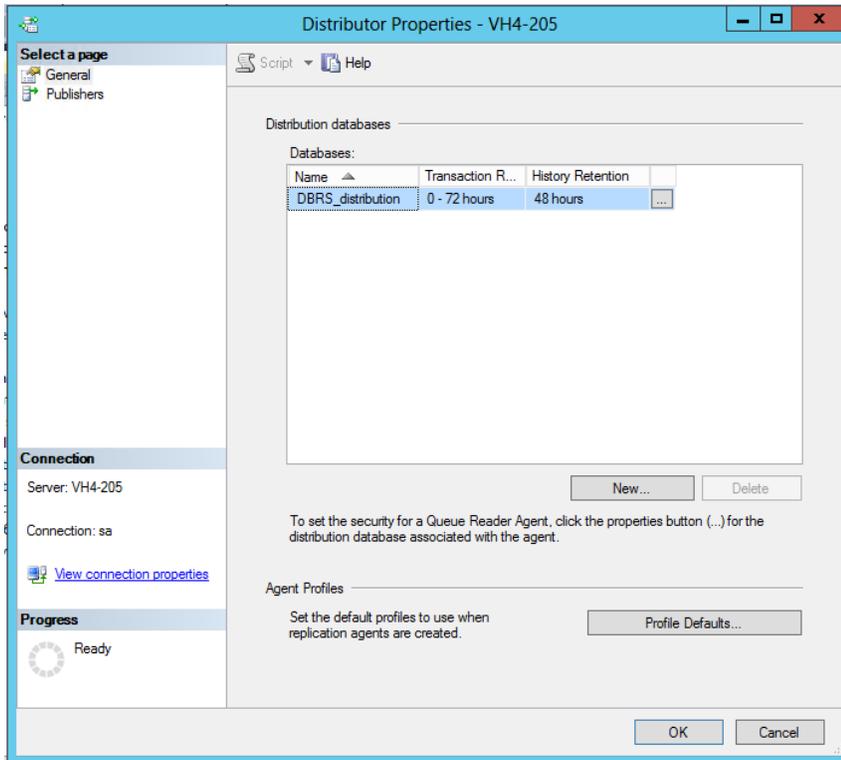
3. Select **Profile Defaults**.
4. Select the **Log Reader Agents** page from the list.
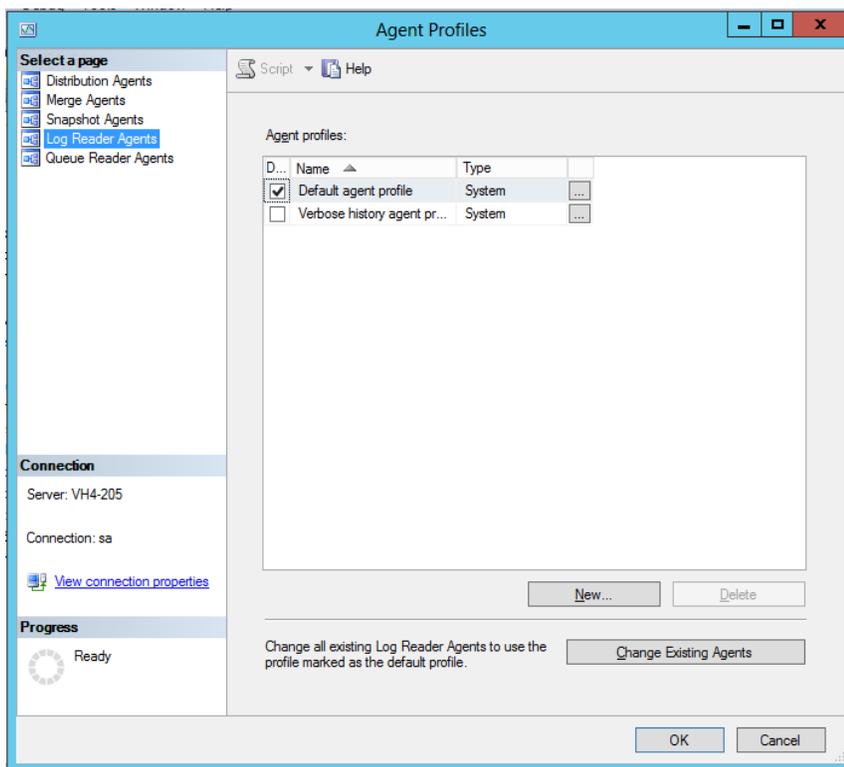
Default agent profile (default for new agents) is selected. To make a copy of this profile and apply a custom setting:

5. Click **New** to open the **New Agent Profile** dialog.
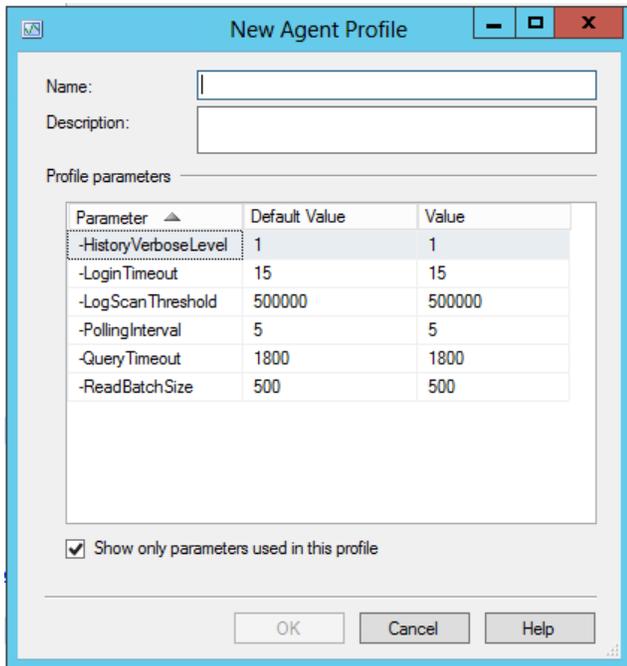
6. Type a name for the new profile.
7. Uncheck **Show only parameters used in this profile**.
8. Scroll to the property called **PublisherFailoverPartner**.

9. Set the name of the secondary replica, VH4-207.

10. Click **OK** to save the profile.

11. In the Agent Profile properties, select the new agent profile as default instead of the **Default agent profile**.
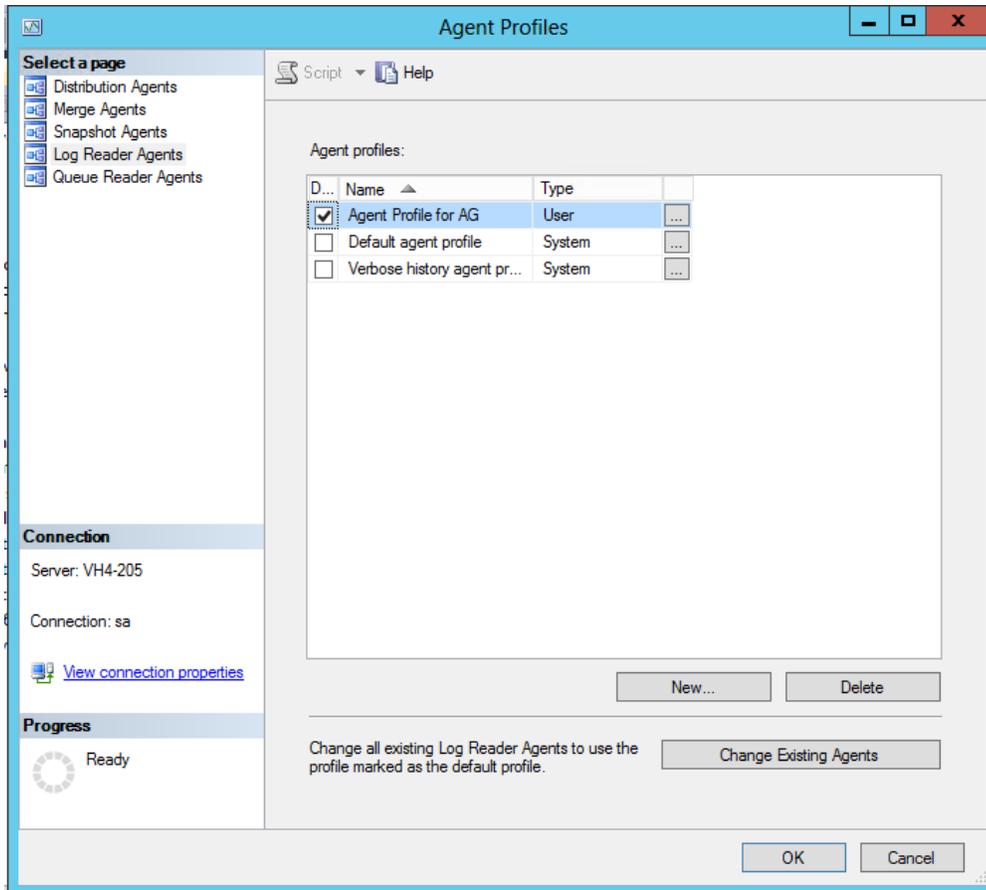
12. Click **OK**.

13. Click **OK** again to exit the Distributor Properties dialog.

NOTE: The steps above for setting the agent profile work for a single secon dary replica. For multiple secondary replicas. it may be possible to define multiple agent profiles, each one using a different PublisherFailoverPartner value, and set it properly as default to each publisher. However, the details are beyond the scope of this document. Please contact HiT Software support at support.hitsw.com for additional information.

5. Add Linked Servers to Secondary Replicas

In the event that a secondary replica transitions to the primary role, it must be configured so that the secondary can take over after a failover.

All possible publishers will connect to the subscriber using a linked server. To create a linked server to the subscriber, VH4-205, open a connection to the secondary replicas and create the linked server to it.

1. Connect to the secondary replica VH4-207.

2. Run the query below:

```
EXEC sys.sp_addlinkedserver @server = 'VH4-205'
```

6. Configure the SQL Server Connection in DBMoto

Now it is time to configure the source connection in DBMoto that points to the SQL Server Availability Group (AG). This is so that the application can connect to the AG and switch dynamically among the cluster replicas.

In the DBMoto Management Center:

1. From the Metadata Explorer **Sources** list, right click and choose **Add New Connection**.
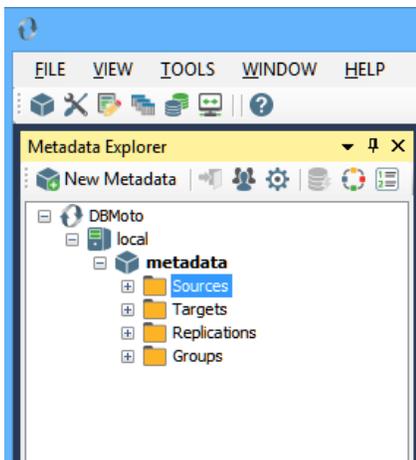
2. In the **Select Provider** screen, type a name and choose **Microsoft SQL Server** as the provider.
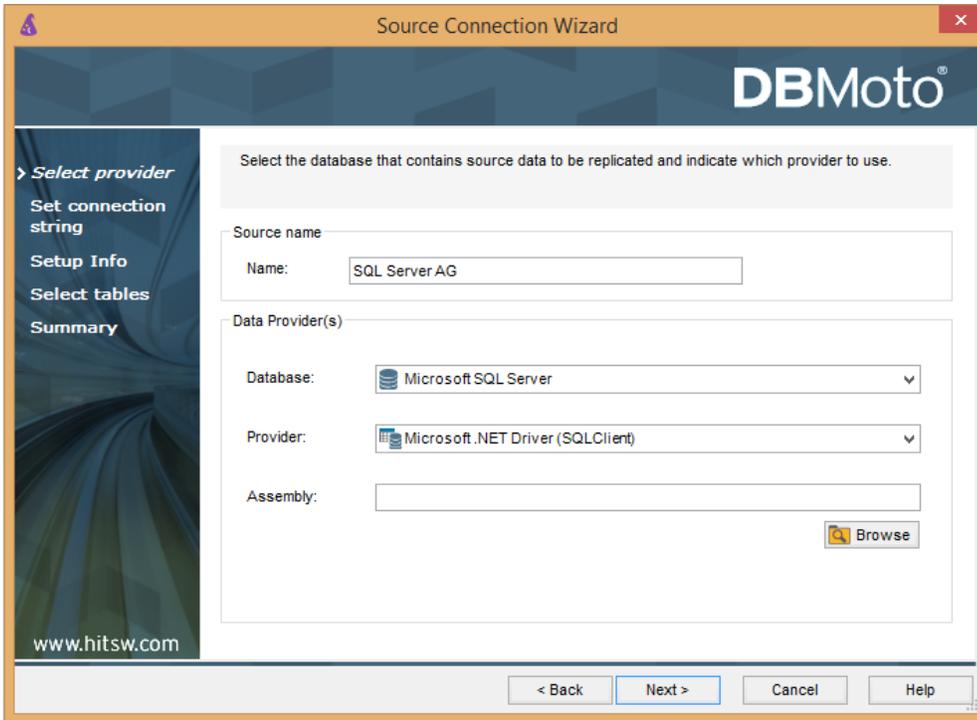
3. Specify connection parameters.
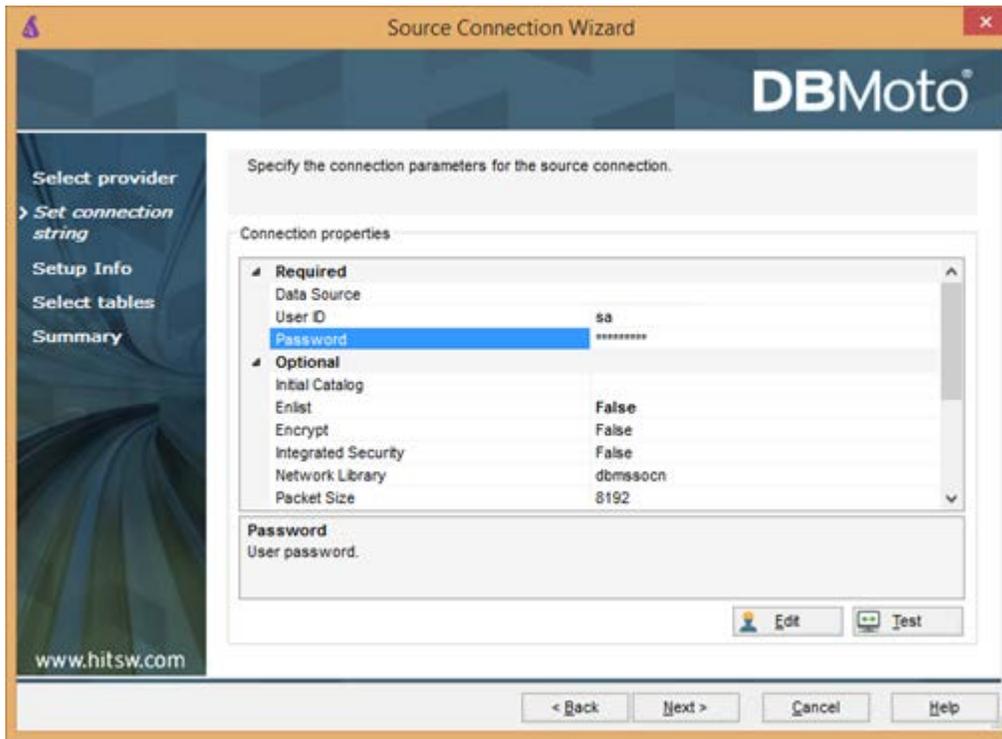
4. In the **Setup Info** screen, select **Log Reader.**

5. Set the IP address that you use to connect to SQL Server from DBMoto to the Availability Group listener IP address.
   Using the listener IP, allows DBMoto to automatically switch from one replica to another once a failover has occurred.

6. Specify the user name and password for the connection. Click **Verify** to check for an existing distributor.
   DBMoto determines that a remote distributor has been established but a login is required to be able to connect to it.

7. The Change Distributor dialog automatically shows the remote distributor name VH4-205, and asks you to specify a user name and password.
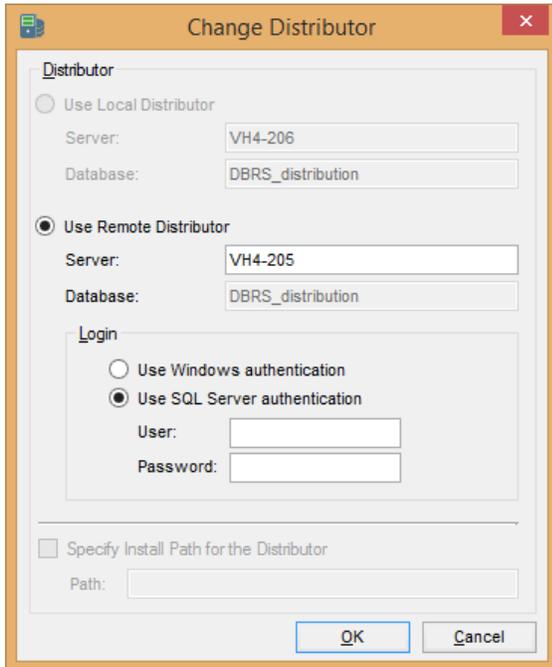
8. Type the user name and password.

9. Click **OK**.

   A message prompt shows that the distributor is correctly set up.

10. Click **Next**.

11. In the **Select Tables** screen, select a few test tables from the AG database (in this case, AGTest1.)
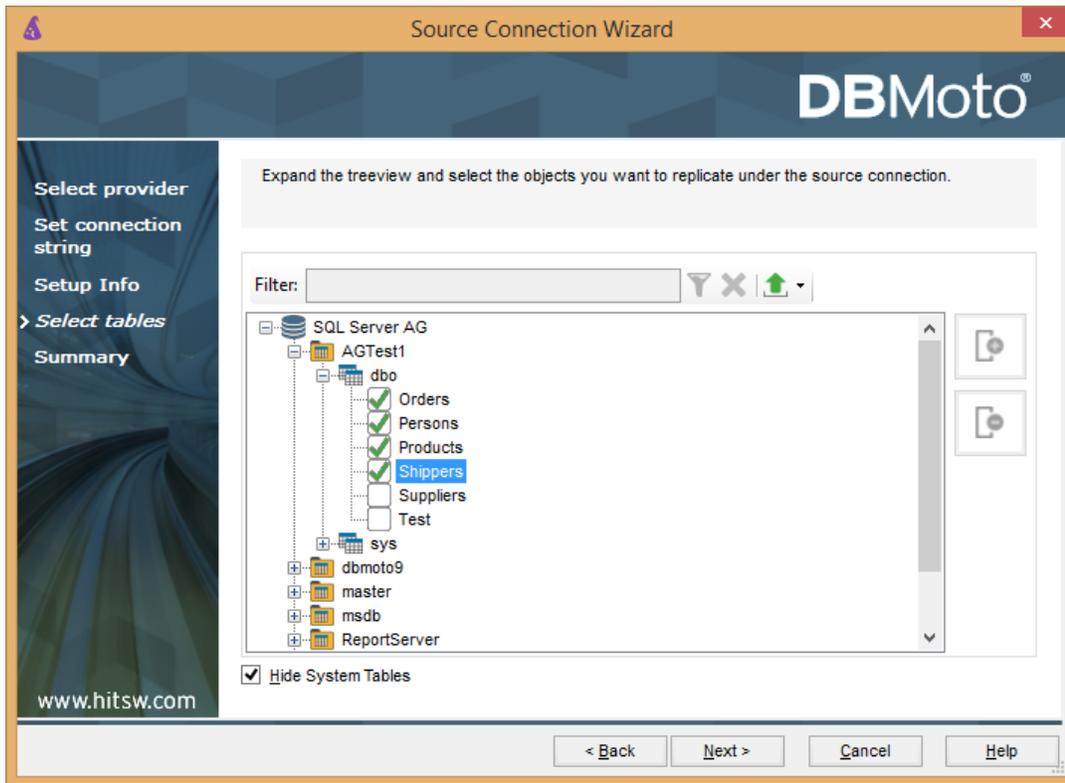
12. Click **Next**.

13. Click **Finish** to create the connection.

7. Add Replications and Create the Publication Database

Add a target connection and create replications in the usual way.

8. Redirect the Publisher to the AG Listener Name

1. In the SQL Server Management Studio, connect to the distributor VH4-205.

2. For **each replica** (VH4-206 and VH4-207), run the stored procedure sp_redirect_publisher to associate the original publisher and the published DB with the AG listener name:

```
USE DBRS_distribution;
GO
EXEC sys.sp_redirect_publisher
     @original_publisher = 'VH4-206',
     @publisher_db = 'AGTest1',
     @redirected_publisher = 'TEST1_AG_Listen';
```

NOTE: Run this stored procedure again, substituting VH4-207 for VH4-206

3. In the distribution database, for **each replica** (VH4-206 and VH4-207), run the stored procedure
    sp_validate_replica_hosts_as_publishers to verify that the replica host is now configured to serve as publisher for the
    published database:

```
USE DBRS_distribution;
GO
DECLARE @redirected_publisher sysname;
EXEC sys.sp_validate_replica_hosts_as_publishers
        @original_publisher = 'VH4-206',
        @publisher_db = 'AGTest1',
        @redirected_publisher = 'TEST1_AG_Listen';
```

NOTE: Run this stored procedure again, substituting VH4-207 for VH4-206

NOTE: This procedure is necessary whenever a new database is added to a replication: the publisher database will have to
be specified in the queries.

Test the Configuration

Let the replications start and run the initial refresh. Run some transactions from the active node (VH4-206) and verify
that they are correctly replicated to the target.

Execute a manual failover to the secondary replica using the Failover wizard in SQL Server. Wait to see the secondary
node become the primary replica. Run some transactions from the active node (VH4-207) and verify that they are
correctly replicated to the target.

Notes

1. The SQL Server Reader Agent has to be running on the Distributor machine only.
2. When creating the distributor, this error can be generated:

```
 Named Pipes Provider: Could not open a connection to SQL Server [53].

 OLE DB provider "SQLNCLI11" for linked server "repl_distributor" returned
 message "Login timeout expired".

 OLE DB provider "SQLNCLI11" for linked server "repl_distributor" returned
 message "A network-related or

 instance-specific error has occurred while establishing a connection to
 SQL Server. Server is not found or not accessible.
```

```
Check if instance name is correct and if SQL Server is configured to allow
remote connections. For more information see SQL Server Books Online.".
```

SOLUTION:

Check if the instance name is correct, by running 'select @@servername'. This name should match the server name you are using in the distributor wizards and in the SQL statements you have run so far, for instance 'VH4-207'. If this is not the case, you will need to reconfigure the name of the local instance by running:

- sp_dropserver 'oldname' (the name found from the select @servername)
- sp_addserver 'newname'   (for instance, VH4-207)
- Restart the SQL Server services

3. When running the failover to the VH4-207, make sure the Log Server Agent is running and has no errors.

Check eventual errors here:

- In the SQL Server Console, click on 'SQL Server Agent'.
- Expand and double click on 'Job Activity Monitor'
- Check the job named to match the database (in the example above 'VH4-206-AGTest1-1')
- In the SQL Server Console, click on 'Replication'.
- Select 'Launch Replication Monitor'.
- Check if all publishers are functioning or show any error icon. In case of errors, expand to find the error items. For example, the following error could occur:

```
The process could not execute 'sp_replcmds' on VH4-207

Status: 0, code: 15517, text: 'Cannot execute as the database principal
because the principal "dbo" does not exist, this type of principal cannot
be impersonated, or you do not have permission.'.
```

SOLUTION:

The error states that the current owner of the job is not a DBO on the Publication Database. Hence the next logical step is to make the job owner the DBO of the database as follows.

- Connect to VH4-207
- Change the owner of the database:

```
USE AGTest1;

sp_changedbowner 'sa'
```