Syniti

Data Replication

Setup Notes for Replications in SQL Always On AG Environment

Version 9.6.0.50

Software Release Date: 9/12/19

Syniti Formerly HIT SOFTWARE

https://support.hitsw.com

Syniti Data Replication

Document History

Version	Author	Date	Reviewer	Date	Comments
1	SUpadhaya	6/19/20	VFarruggio	6/23/20	

Syniti Data Replication

Table of Contents

Overview 1
Environment 1
Configuring the Publisher DB in Availability Group 1
Configuring the Distribution DB in Availability Group 9
Additional Workflow 14
Configuration of replication in SDR 17

Syniti Data Replication

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 (AG) and the distribution database is also part of an availability group. It uses the environment described in next section as an example.

SQL Server 2017 CU6 and SQL Server 2016 SP2-CU3 introduced support for replication distribution database in an AG. When the publisher DB is part of an AG and the distribution DB is a remote stand-alone instance, please refer to the document named **MSSQLServer-AlwaysOn-DBMoto9.pdf** in the Help Center to correctly configure publisher and distributor.

Here is a brief overview as it relates to setting up Always On Availability Groups with Syniti DR. 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 databases 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 Syniti Data Replicaton is described below.

Distribution Database	Publisher Database
Cluster Name: WS19DISTRCLSTR1	Cluster Name: WS19ALWONCLSTR1
Nodes: WS19-ALWON-MAC1 (Primary Replica) WS19-ALWON-MAC2 (Secondary Replica)	Nodes: WS19-ALWON-SQL1 (Primary Replica) WS19-ALWON-SQL2 (Secondary Replica)
AG Name: Distributor_AG	AG Name: SyniAG
AG Listener: DistAGListener	AG Listener: SyniAGListener
DB Name: DBRS_distribution	DB Name: syni

Configuring the Publisher DB in Availability Group

The Publisher DB can be any source database that contains the data which needs to be replicated. This document uses **syni** as an example of a Publisher database. The Availability Group for the Publisher DB can be created via the Availability Group Wizard in SQL Server Management Studio.

Before starting with the wizard, make sure endpoints for database mirroring and login accounts on the server instances of primary and replicas are created.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

1. Create endpoints on primary and secondary replicas for database mirroring.

```
-- Run below script on WS19-ALWON-SQL1 and WS19-ALWON-SQL2
CREATE ENDPOINT [Hadr_endpoint]
    STATE=STARTED
    AS TCP (LISTENER_PORT = 5022, LISTENER_IP = ALL)
    FOR DATA_MIRRORING (ROLE = ALL, AUTHENTICATION = WINDOWS NEGOTIATE
, ENCRYPTION = REQUIRED ALGORITHM AES)
Go
```

2. Create a login account in the primary replica's server instance for the secondary replica to connect. Then, that login must be granted CONNECT permissions to connect to the database mirroring endpoint of that server instance.

```
-- Run below script on WS19-ALWON-SQL1(primary replica)
USE [master]
Go

CREATE LOGIN [HITSW\WS19-ALWON-SQL2$] FROM WINDOWS WITH
DEFAULT_DATABASE=[master], DEFAULT_LANGUAGE=[us_english]
Go

GRANT CONNECT on ENDPOINT::[Hadr_Endpoint] TO [HITSW\WS19-ALWON-SQL2$];
```

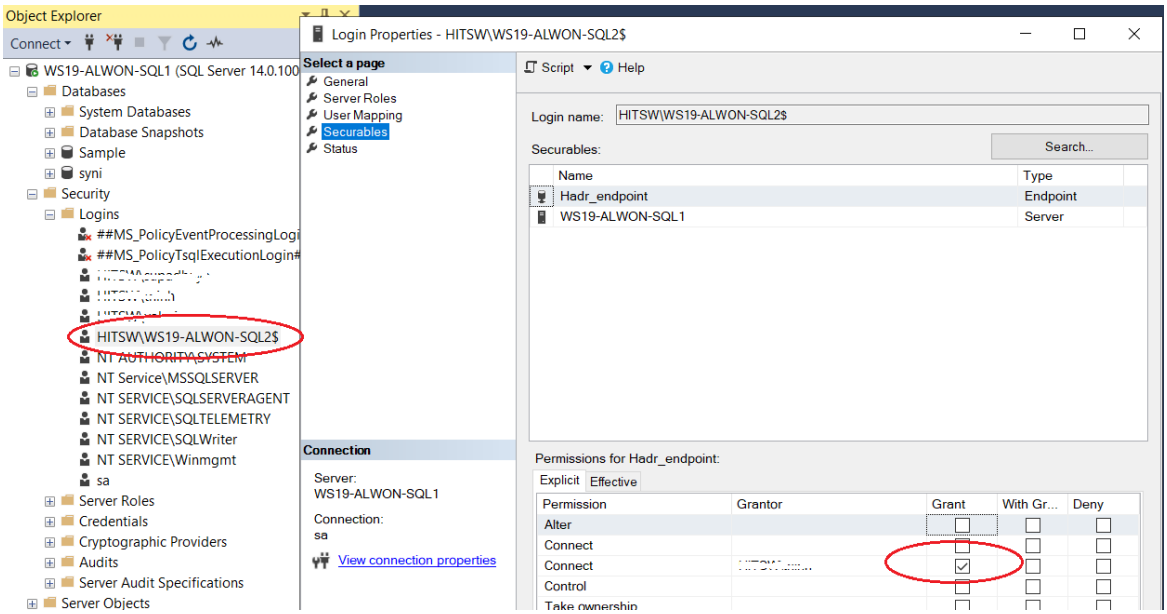

3. Create a login account in the secondary replica's server instance for the primary replica to connect. Then, that login must be granted CONNECT permissions to connect to the database mirroring endpoint of that server instance.

```
-- Run below script on WS19-ALWON-SQL2(secondary replica)
USE [master]
```

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

```
Go

CREATE LOGIN [HITSW\WS19-ALWON-SQL1$] FROM WINDOWS WITH
DEFAULT_DATABASE=[master], DEFAULT_LANGUAGE=[us_english]
Go

GRANT CONNECT on ENDPOINT::[Hadr_Endpoint] TO [HITSW\WS19-ALWON-SQL1$];
```

4. The Publisher DB **syni** is only present in the primary replica, with recovery set to “Full”.

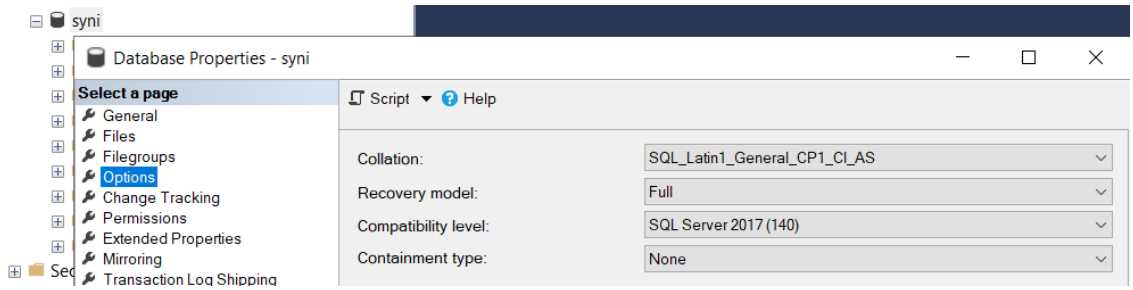

5. Start the wizard to create the availability group.

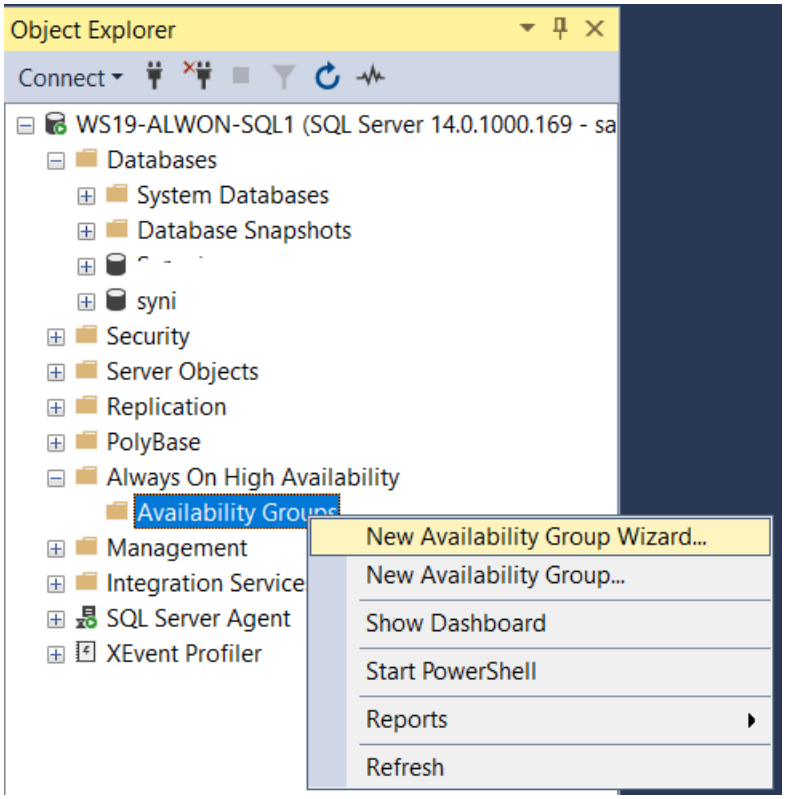

6. Provide the AG name and click **Next**.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

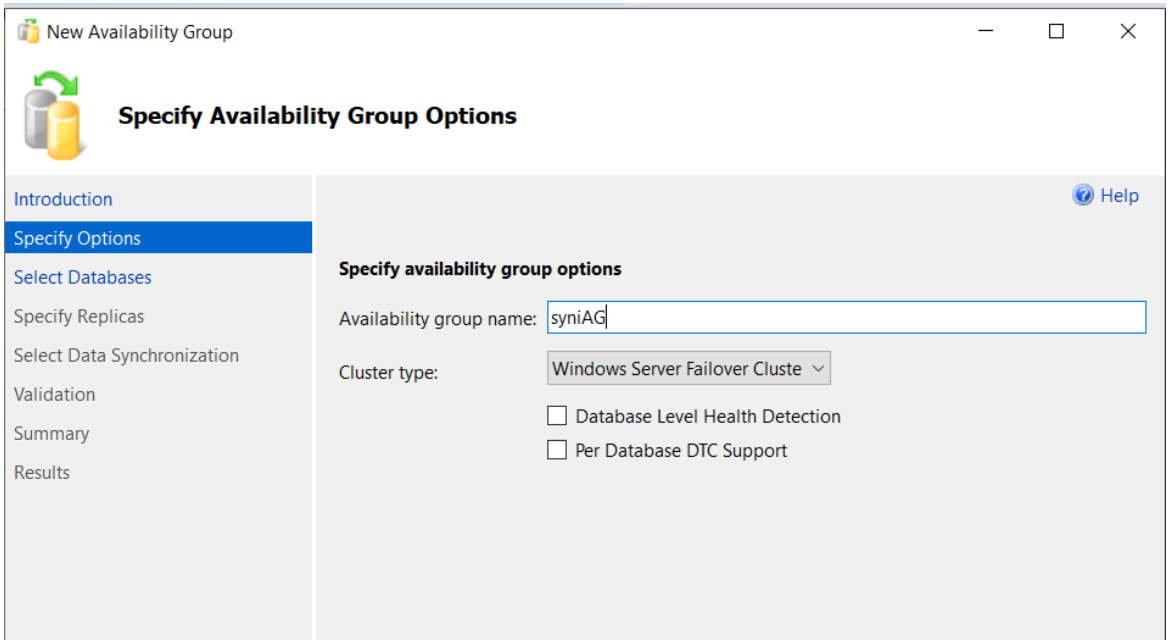

7. Select database "**syni**" from the options and click **Next**.

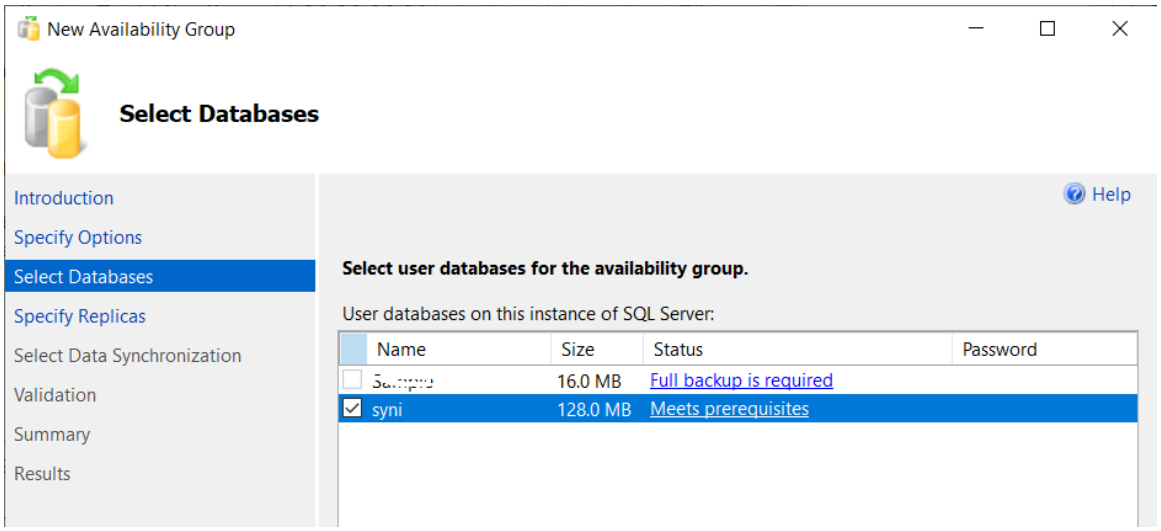

8. Set readable secondary to **yes** and availability mode to **Synchronous commit** and then click **Add Replica**.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

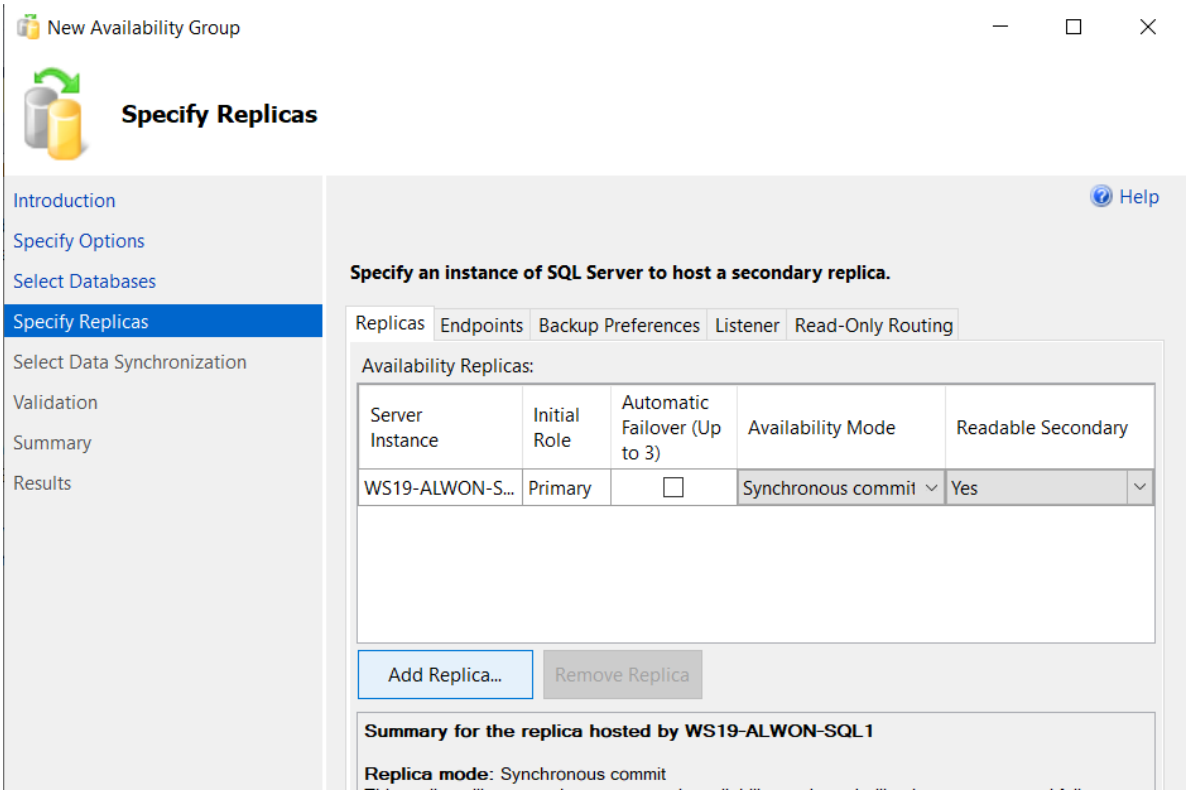

9. Add the secondary replica and set properties as shown in the screenshot below.

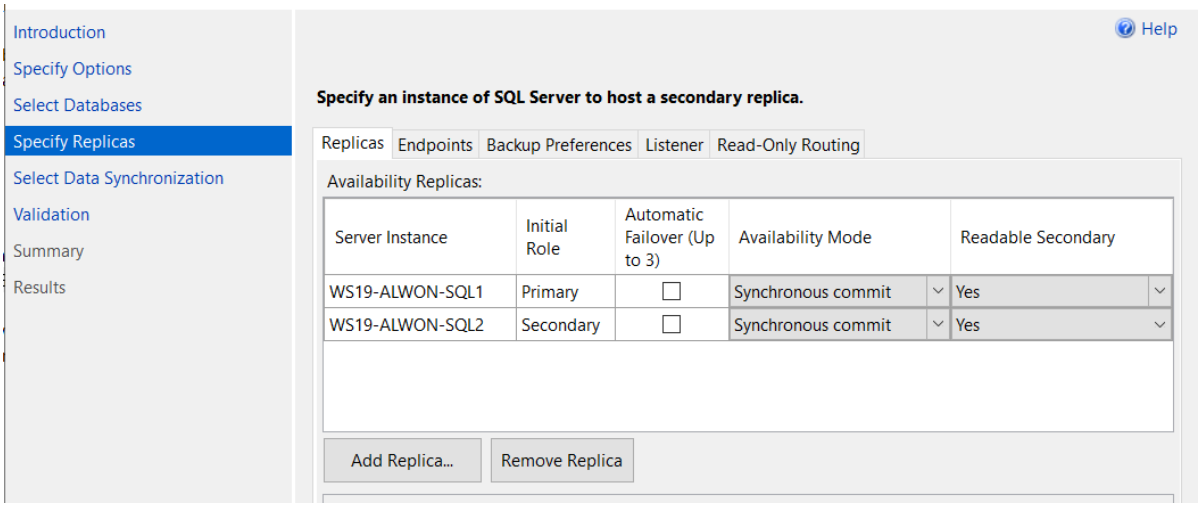

10. Select the **Automatic seeding** option in the “Select Data Synchronization” step.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

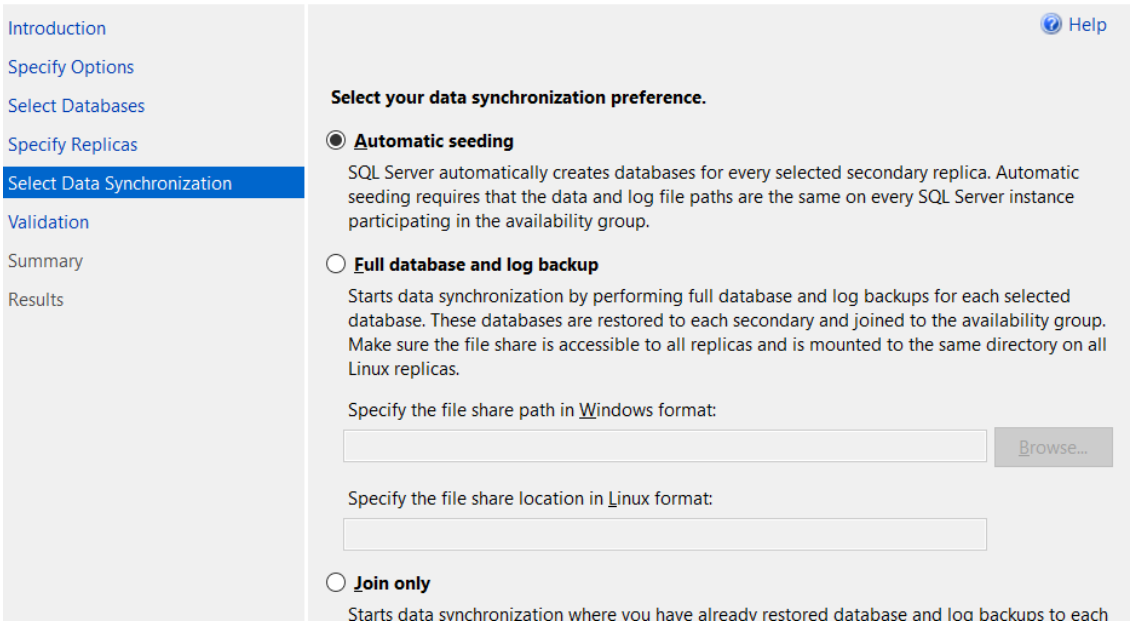

11. Click **Next** on the validation step.

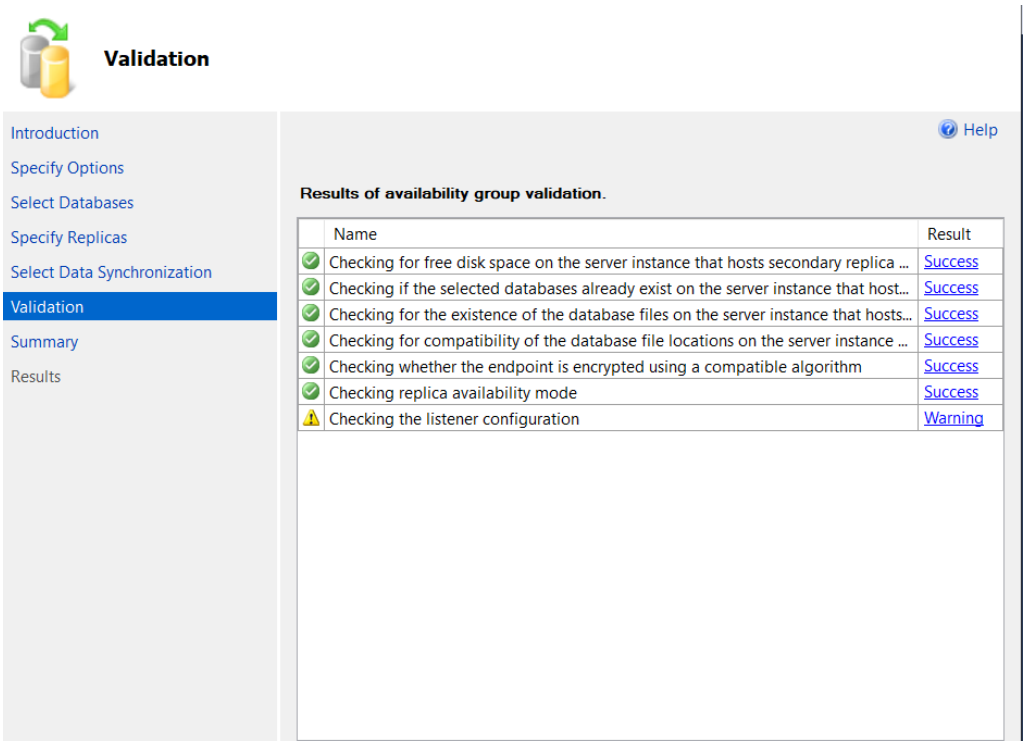

12. Click **Finish** to complete the wizard successfully.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

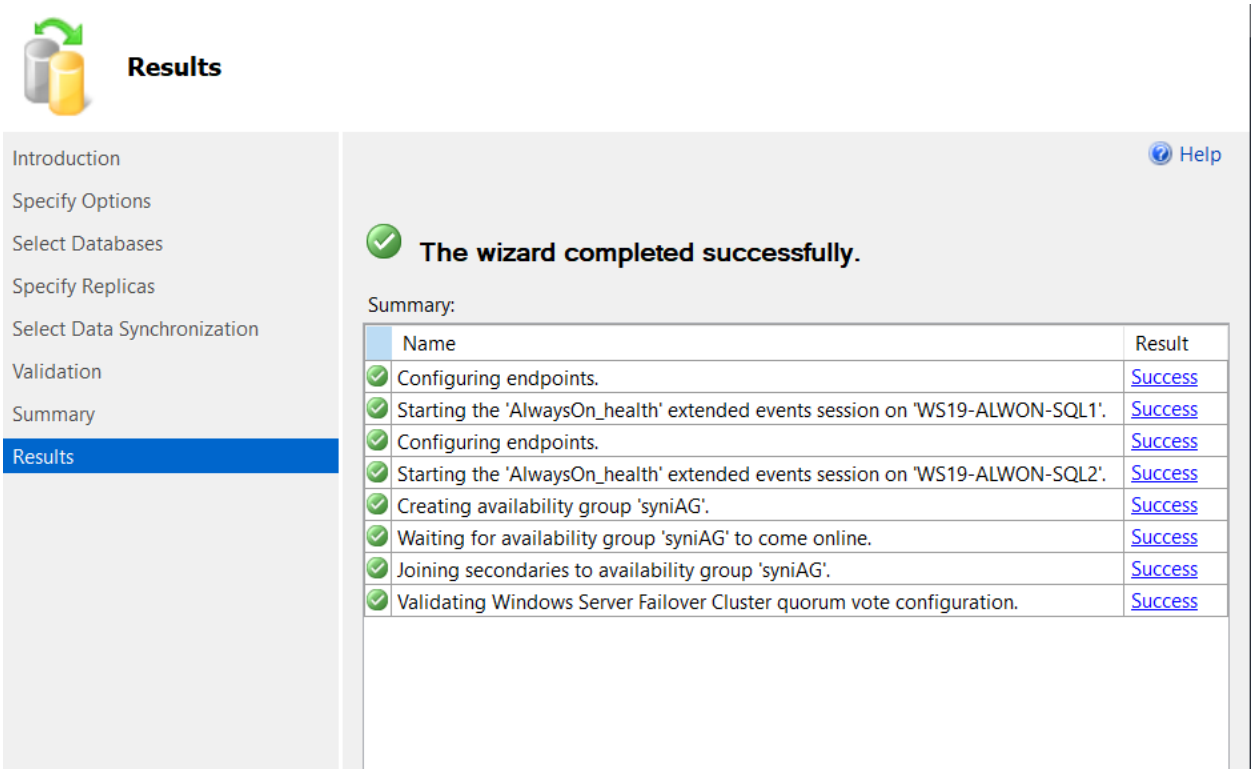

13. Add the listener to the AG.

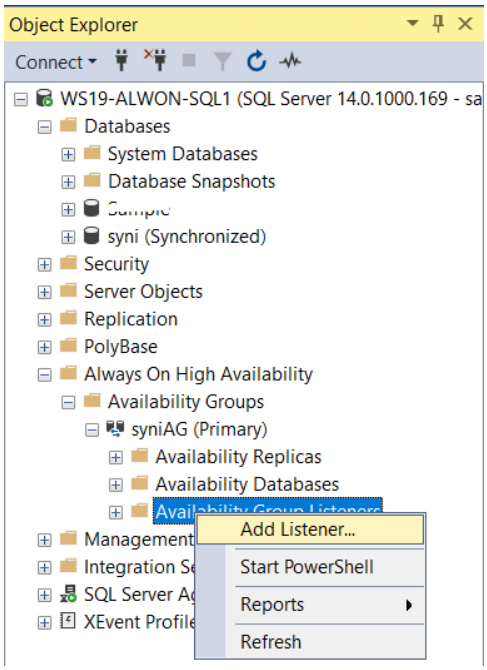

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

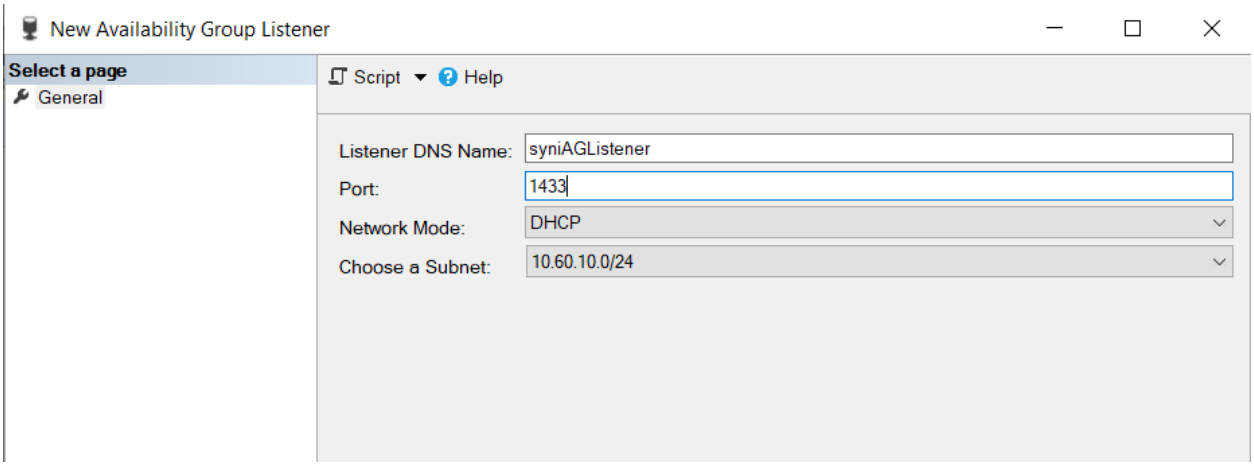

14. The Availability Group's dashboard should look like the screenshot below.

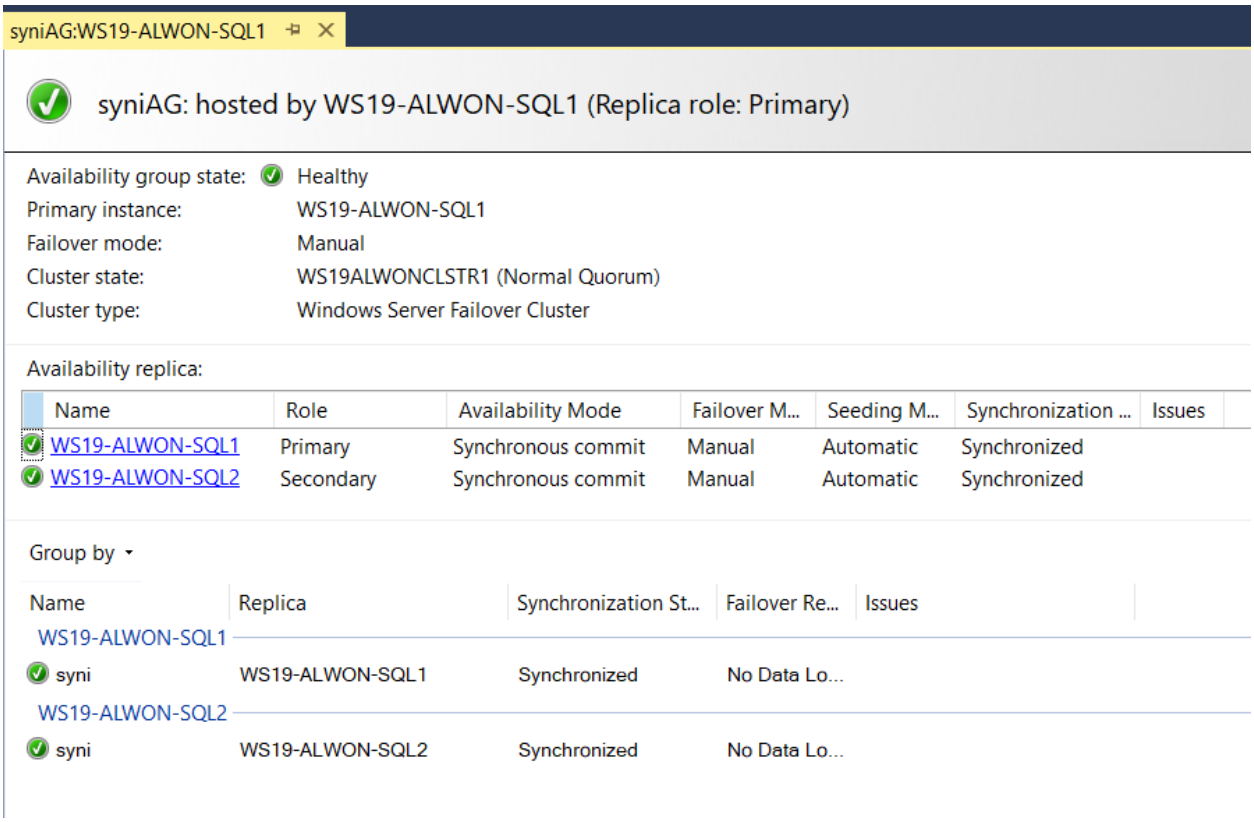

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

Configuring the Distribution DB in Availability Group

The distribution database AG needs to have a listener. When the Publisher adds the Distributor, it uses the listener name as the Distributor name. The log reader agent jobs created on the distribution server get created on all secondary replicas of the AG for Distribution DB. Please refer to below document for detailed information about setting up distribution databases in an AG.

https://docs.microsoft.com/en-us/sql/relational-databases/replication/configure-distribution-availability-group?view=sql-server-ver15

1. Create endpoints on primary and secondary replicas for database mirroring.

```
-- Run below script on WS19-ALWON-MAC1 and WS19-ALWON-MAC2
CREATE ENDPOINT [Hadr_endpoint]
    STATE=STARTED
    AS TCP (LISTENER_PORT = 5022, LISTENER_IP = ALL)
    FOR DATA_MIRRORING (ROLE = ALL, AUTHENTICATION = WINDOWS NEGOTIATE
, ENCRYPTION = REQUIRED ALGORITHM AES)
```

2. Create a login account in the primary replica's server instance for the secondary replica to connect. Then, that login must be granted CONNECT permissions to connect to the database mirroring endpoint of that server instance.

```
-- Run below script on WS19-ALWON-MAC1(primary replica)
USE [master]
Go

CREATE LOGIN [HITSW\WS19-ALWON-MAC2$] FROM WINDOWS WITH
DEFAULT_DATABASE=[master], DEFAULT_LANGUAGE=[us_english]
Go

GRANT CONNECT on ENDPOINT::[Hadr_Endpoint] TO [HITSW\WS19-ALWON-MAC2$];
```

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

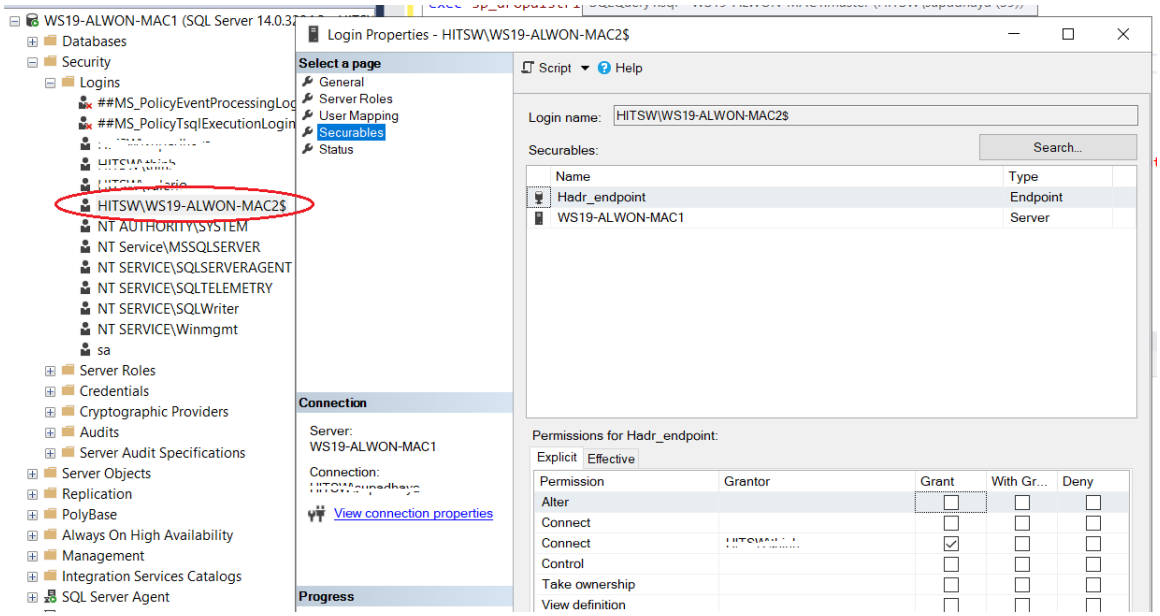

3. Run the script below on the secondary replica to create a login account for the primary replica to connect.

```
-- Run below script on WS19-ALWON-MAC2
USE [master]
Go

CREATE LOGIN [HITSW\WS19-ALWON-MAC1$] FROM WINDOWS WITH
DEFAULT_DATABASE=[master], DEFAULT_LANGUAGE=[us_english]
Go

GRANT CONNECT on ENDPOINT::[Hadr_Endpoint] TO [HITSW\WS19-ALWON-MAC1$];
```

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

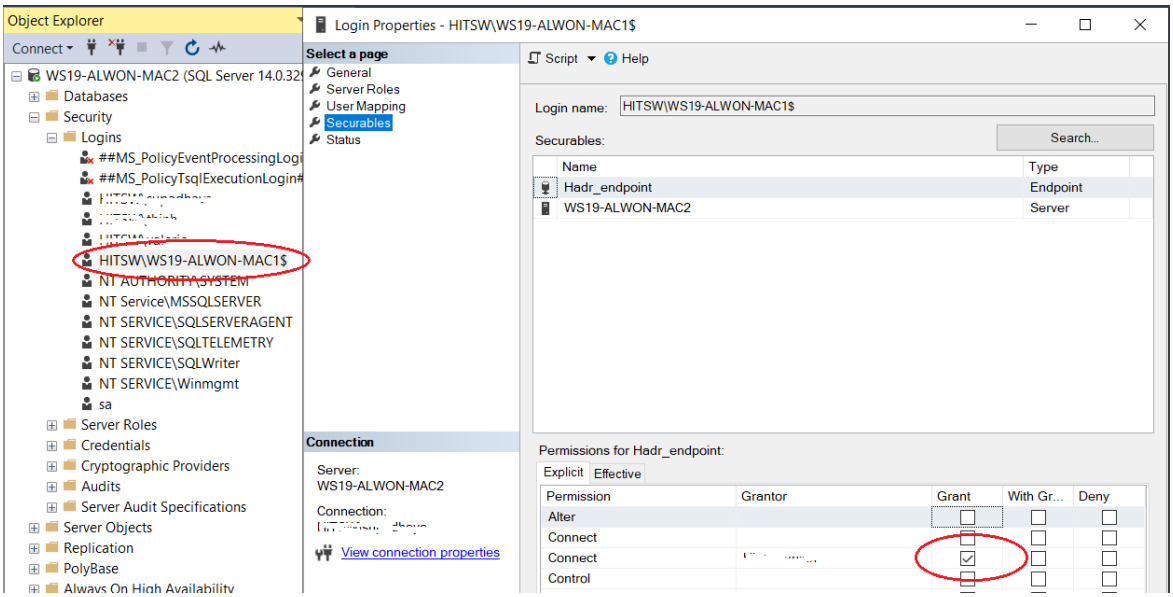

4. Connect to the system WS19-ALWON-MAC1 (primary replica) from MS SQL Server Management Studio and run the script below.

```
Use master
Go
-- Add distributor.
--Specify the password for distributor_admin through the @password. The
@password should be identical across 'WS19-ALWON-MAC1' and 'WS19-ALWON-MAC2'

exec sp_adddistributor @distributor = 'WS19-ALWON-MAC1', @password =
'Password'
Go

-- Create distribution DB 'DBRS_distribution
exec sp_adddistributiondb @database = 'DBRS_distribution', @data_file =
'DBRS_distribution_Data', @data_file_size = 5, @log_file =
'DBRS_distribution_Log', @log_file_size = 2, @max_distretention = 72,
@security_mode = 1
Go

-- Set Recovery to Full
Alter Database [DBRS_distribution] Set Recovery Full
Go

-- Take DB backup
Backup Database [DBRS_distribution] to Disk = 'Nul'
Go
```

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

5. Create the Availability Group **Distributor_AG**

```
-- Create availability group for 'DBRS_distribution'.
USE [master]
Go
CREATE AVAILABILITY GROUP [Distributor_AG]
  FOR DATABASE DBRS_distribution
  REPLICA ON N'WS19-ALWON-MAC1'
WITH (
  ENDPOINT_URL = N'TCP://WS19-ALWON-MAC1.hitsw.com:5022',
  FAILOVER_MODE = AUTOMATIC,
  AVAILABILITY_MODE = SYNCHRONOUS_COMMIT,
SECONDARY_ROLE(ALLOW_CONNECTIONS = ALL),
SESSION_TIMEOUT = 10
),
  N'WS19-ALWON-MAC2' WITH (
    ENDPOINT_URL = N'TCP://WS19-ALWON-MAC2.hitsw.com:5022',
    FAILOVER_MODE = AUTOMATIC,
AVAILABILITY_MODE = SYNCHRONOUS_COMMIT,
    SEEDING_MODE = AUTOMATIC,
SECONDARY_ROLE(ALLOW_CONNECTIONS = ALL),
SESSION_TIMEOUT = 10);
 Go
```

Add a listener to the AG from Object Explorer. Provide the default port 1433 when creating the listener.

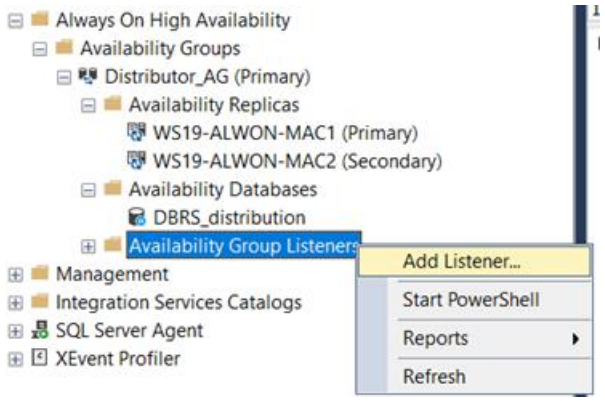

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

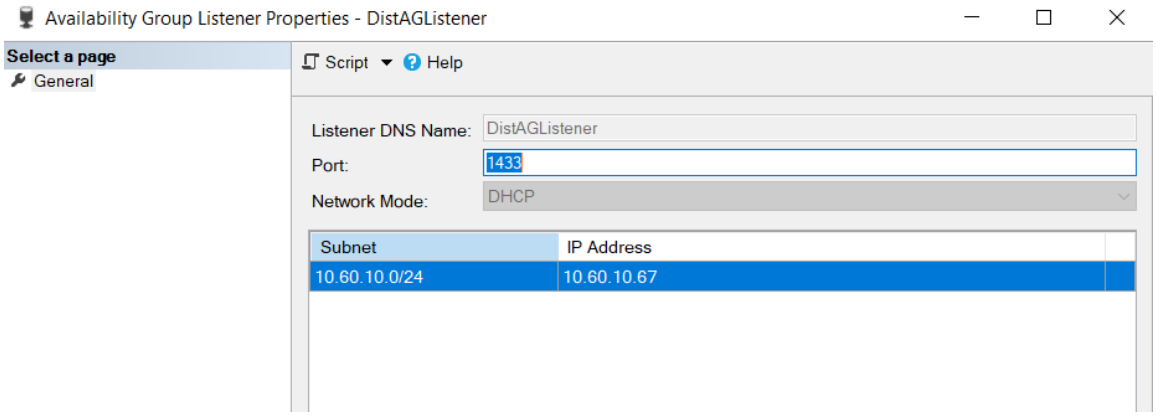

6. Run the script below on the secondary replica (system: WS19-ALWON-MAC2)

```
USE master
Go
--Specify the same password as provided when running the sp on 'WS19-ALWON-
MAC1'
exec sp_adddistributor @distributor = 'WS19-ALWON-MAC2', @password =
'Password'
Go
ALTER AVAILABILITY GROUP [Distributor_AG] JOIN
GO

ALTER AVAILABILITY GROUP [Distributor_AG] GRANT CREATE ANY DATABASE
Go

exec sp_adddistributiondb @database = 'DBRS_distribution', @security_mode = 1
```

7. Run the script below to add WS19-ALWON-SQL1 and WS19-ALWON-SQL2 as publisher on WS19-ALWON-MAC1 and WS19-ALWON-MAC2

```
--Run below script on WS19-ALWON-MAC1 and WS19-ALWON-MAC2
USE master
 Go
EXEC sp_addDistPublisher @publisher = 'WS19-ALWON-SQL1', @distribution_db =
'DBRS_distribution',
    @working_directory =
'\\hitsnap03\Public\SHARED\DistributionDB_WorkingFolder\'
Go
    EXEC sp_addDistPublisher @publisher = 'WS19-ALWON-SQL2', @distribution_db
= 'DBRS_distribution',
    @working_directory =
'\\hitsnap03\Public\SHARED\DistributionDB_WorkingFolder\'
```

The value of @working_directory should be a network path independent of WS19-ALWON-MAC1 and WS19-ALWON-MAC2.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

8. All the nodes in the distribution database AG need to use the same domain account to run SQL Server Agent, and this domain account needs to have the same privilege on each node.

Additional Workflow

After AGs for Distributor and Publisher are configured properly, follow the steps below.

1. To add the distribution database AG listener as the Distributor, on Publisher DB replicas, run:

```
--Run below script on WS19-ALWON-SQL1 and WS19-ALWON-SQL2

USE master
Go
EXEC sp_adddistributor @distributor = 'DistAGListener', @password =
'Password'
```

   The value of @password should be the one that was specified when Distributors were configured.

2. After adding the Distributor on the Publisher, enable replication for the publisher database **syni**.

```
-- run only on server instance used as primary replica (WS19-ALWON-SQL1)
Use master
Go
EXEC sys.sp_replicationdboption
@dbname = 'syni',
@optname = 'publish',
@value = 'true';
```

3. Follow steps in **Configuration of Replication in SDR** to create a replication.

4. After the publication has been created on the Publisher database, run the stored procedure sp_redirect_publisher to associate the original Publisher and the published DB with the AG listener name. Run the script below on the primary replica of Distribution DB.

```
-- Run on Primary replica (WS19-ALWON-MAC1), add Publisher's listener
USE DBRS_distribution
GO

EXEC sys.sp_redirect_publisher
@original_publisher = 'WS19-ALWON-SQL1',
@publisher_db = 'syni',
@redirected_publisher = 'SyniAGListener'
GO
```

5. 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.

```
-- Run on Primary replica (WS19-ALWON-MAC1)
USE DBRS_distribution;
```

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

```
GO

DECLARE @redirected_publisher sysname;
EXEC sys.sp_validate_replica_hosts_as_publishers
@original_publisher = 'WS19-ALWON-SQL1',
@publisher_db = 'syni',
@redirected_publisher = @redirected_publisher output;

PRINT @redirected_publisher --Output: SyniAGListener
```

6. Using the Replication Monitor, register the Publisher's listener **syniAGListener** on the Distributor's Primary replica.

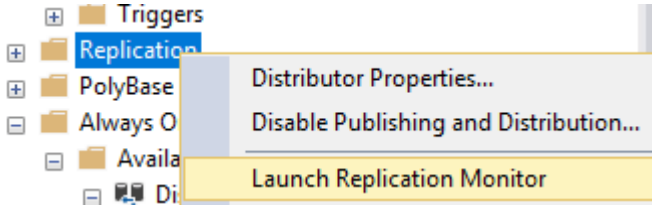

7. Click on the **Add Publisher** link.

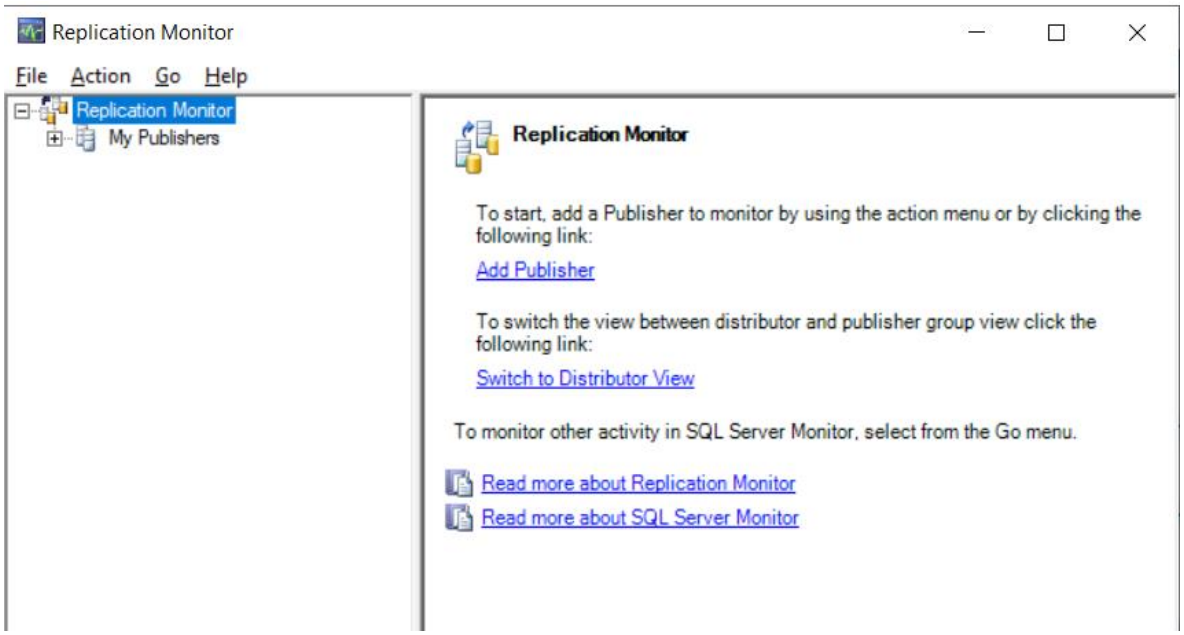

8. Click on **Add SQL Server Publisher** as shown.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

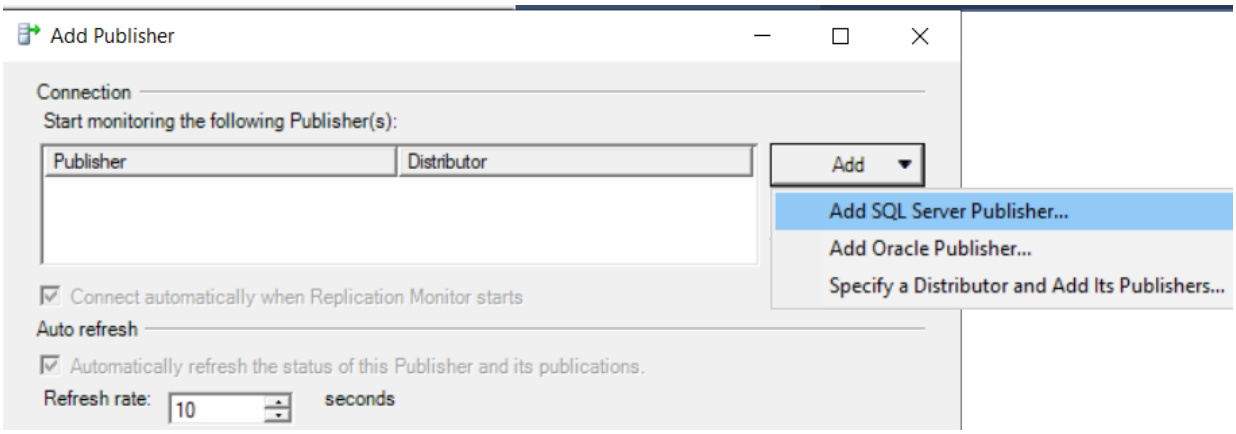

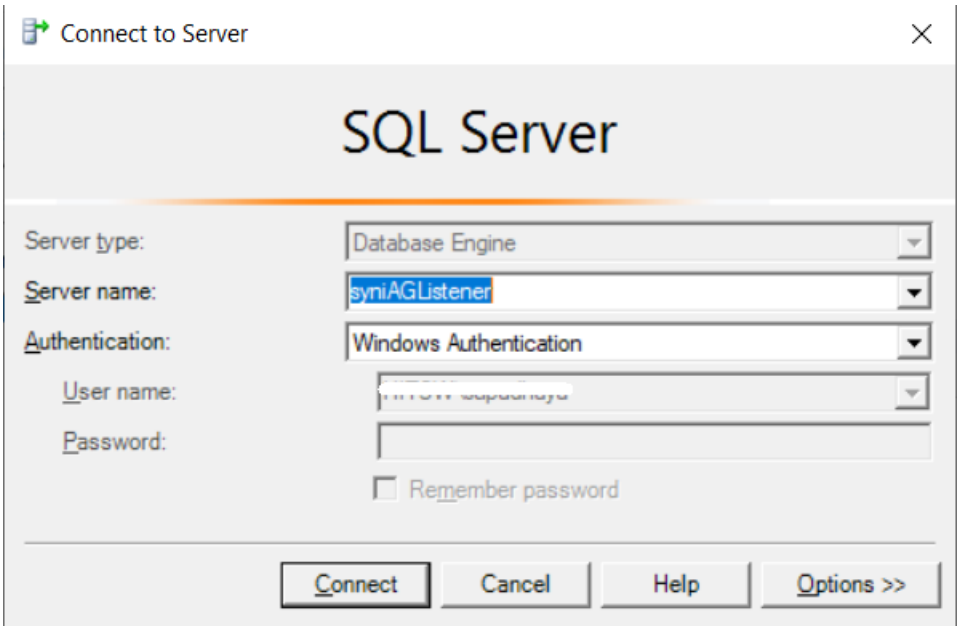

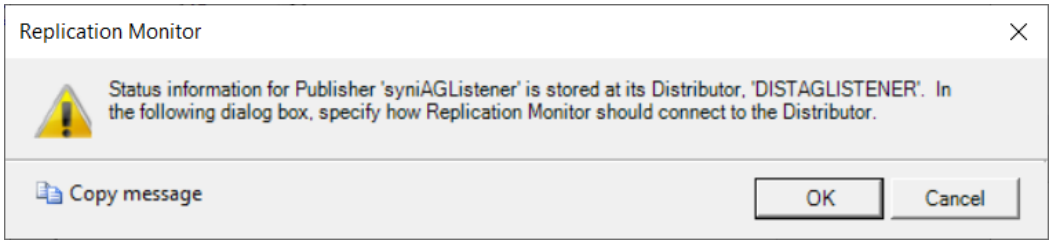

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

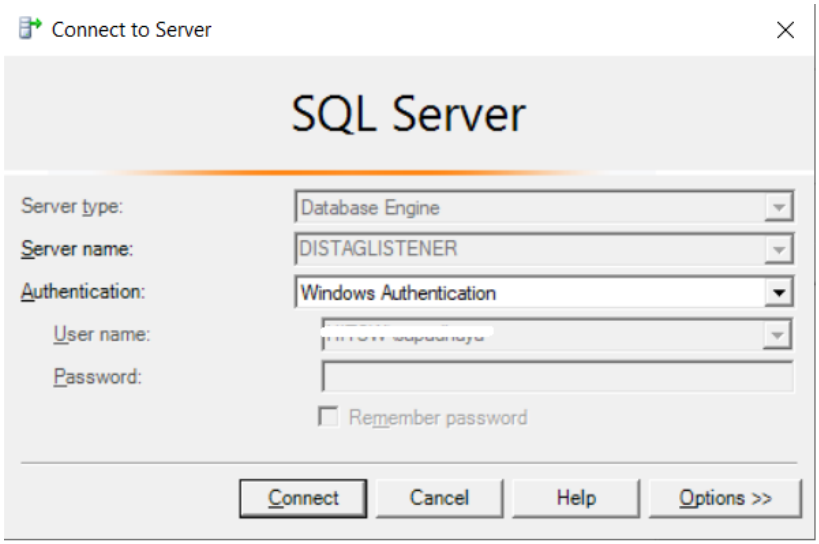

9. Connect to the Distributor's listener to retrieve publisher settings.

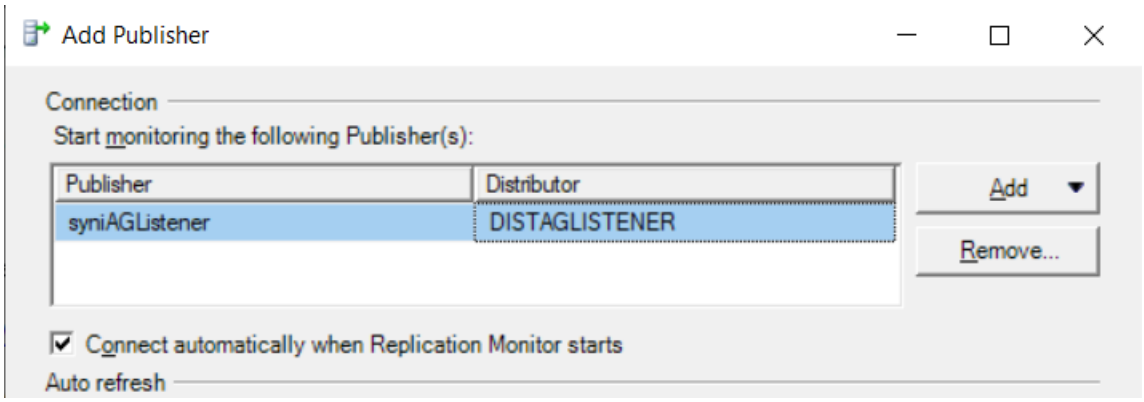

Configuration of Replication in SDR

1. Create a source connection to the Publisher DB from SDR. The Publisher's AG Listener name should be used as **Data Source** in the connection string (as shown in the screenshot), instead of one of the node's instance names.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

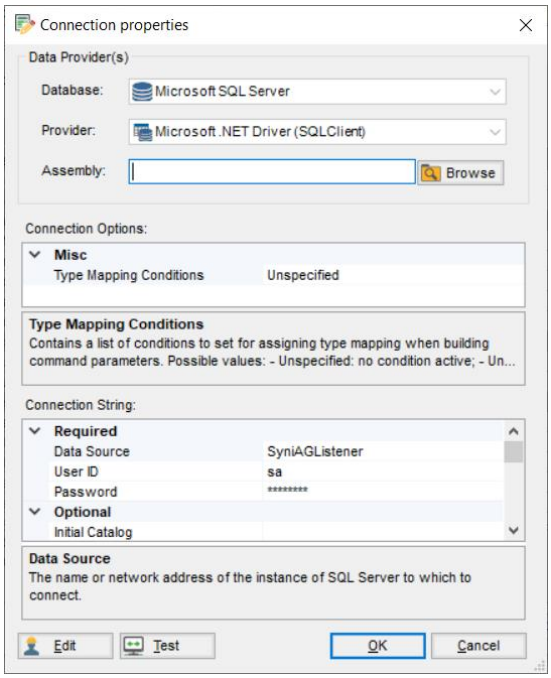

2. Enable a transactional replication using the Enable Transactional Replication Setup wizard in SDR.

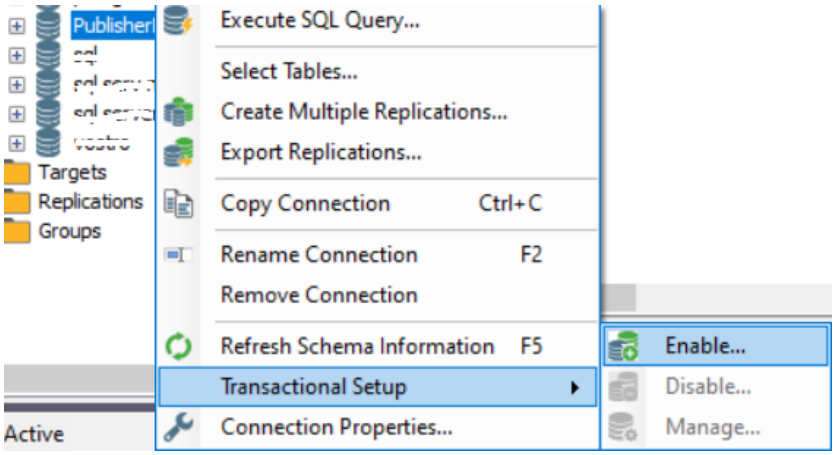

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

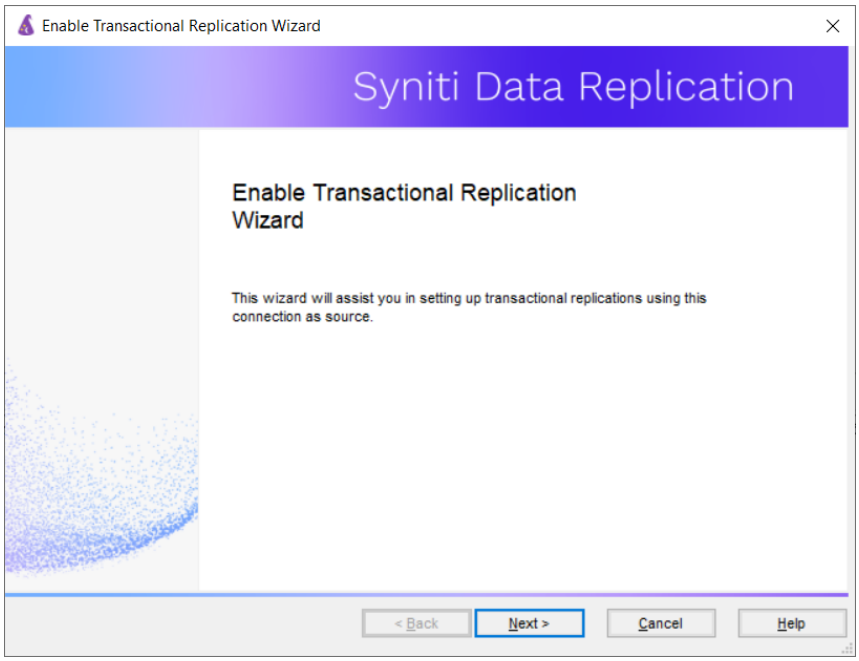

3. Choose **Log Reader** as the Log Type.

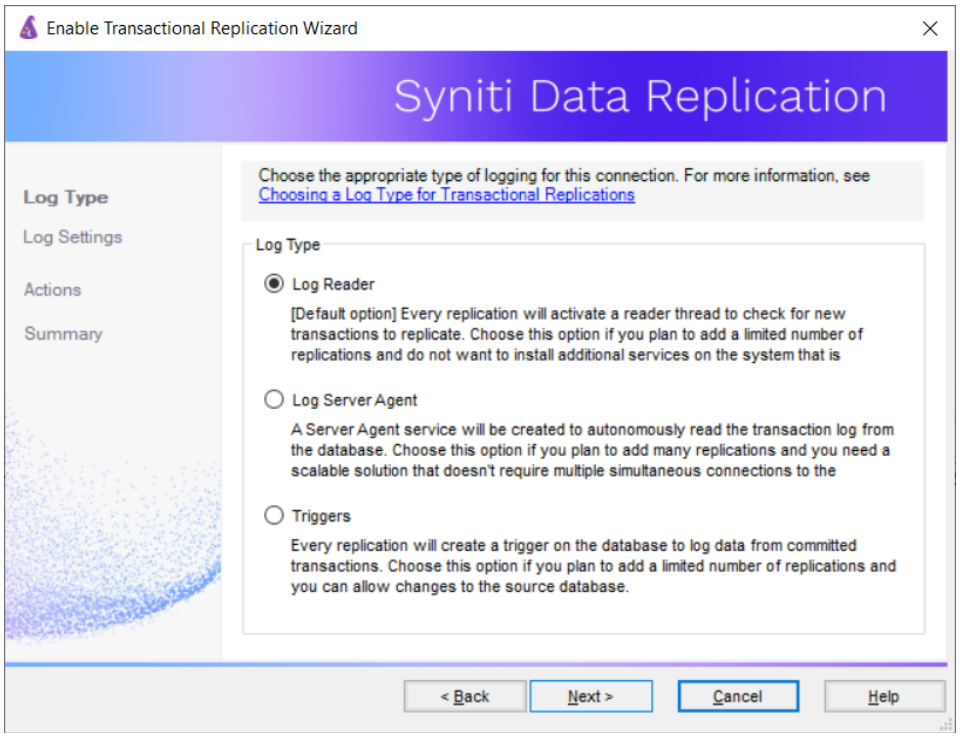

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

4. Provide system admin credentials and click on the highlighted button to make sure the remote distributor is selected.

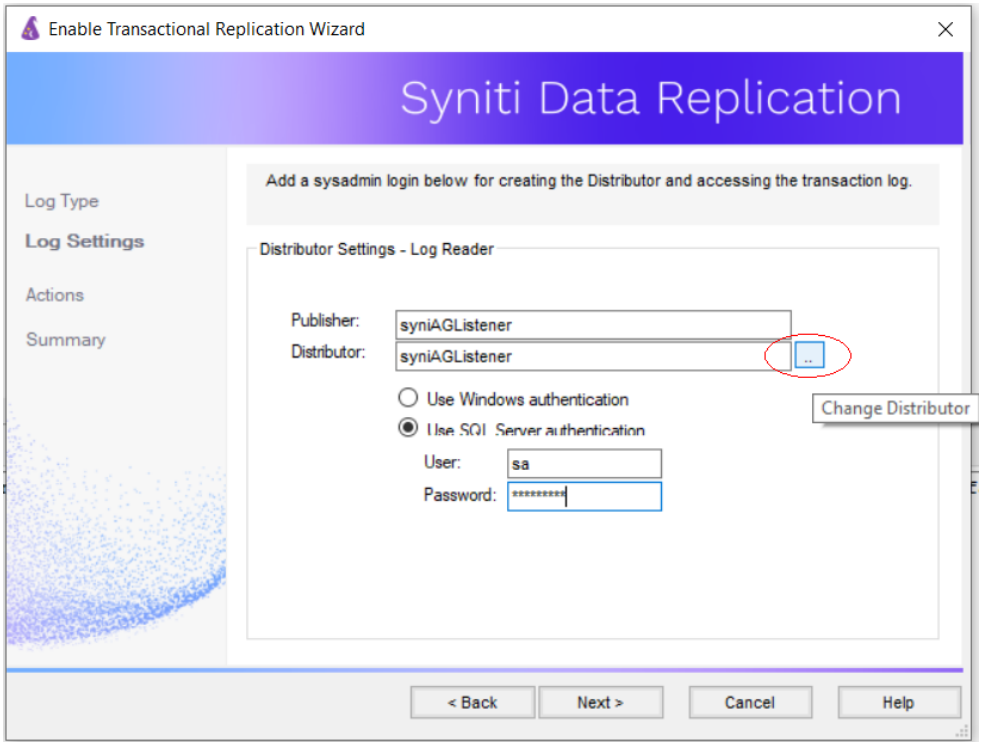

5. Provide login credentials for the remote distributor and click **OK**.

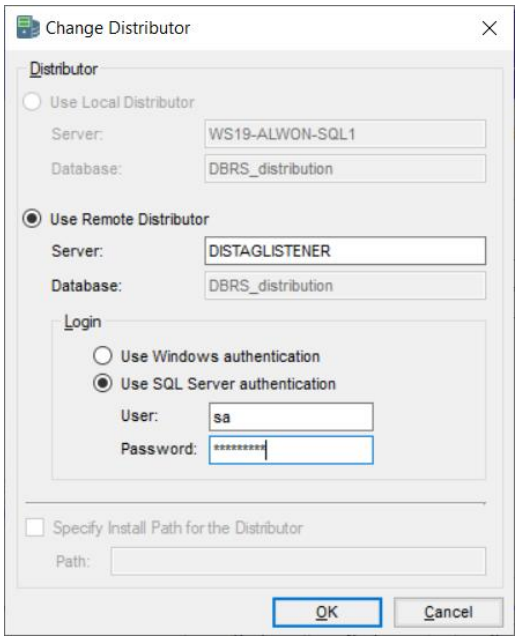

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.

6. Add a target connection and create replications in the usual way. This will create publications on the Publisher DB.

Copyright © 2020 HiT Software, Inc d/b/a Syniti and/or its affiliates. All rights reserved. This document contains confidential and proprietary information and reproduction is prohibited unless authorized by Syniti. Names appearing within the product manuals may be trademarks of their respective owners.