

# DEPLOYING IBM INFOSPHERE DATA REPLICATION AGAINST A MICROSOFT SQL SERVER DATABASE PARTICIPATING IN AN ALWAYSON AVAILABILITY GROUP

## TARGET APPLY

# Table of Contents

1	Deploying IBM IIDR in Microsoft SQL Server AlwaysOn availability group .....	3
2	Local apply.....	3
2.1	Overview .....	3
2.2	Add symmetrical shared clustered disk reachable from all AlwaysOn nodes .....	5
2.3	Define a windows domain user that can be used to run the IBM IIDR service .....	8
2.4	IIDR Database authentication user .....	8
2.5	Deploy IBM IIDR on the shared clustered disk.....	9
2.6	Create IIDR service on secondary synchronous replicas.....	11
3	Remote apply .....	16
3.1	Overview .....	16
3.2	Deploy IBM IIDR in remote apply topology.....	17
4	Configure the new Datastore in Management Console .....	20
5	Test the deployment .....	20
6	Persistent subscriptions .....	20
7	Considerations when transitioning the database to an AlwaysOn availability group.....	21
7.1	Using a new IIDR installation and instance .....	21
7.2	Reusing the existing installation and instance in remote apply topology .....	21
8	Considerations when transitioning the database out of an AlwaysOn availability group .....	22
8.1	Using existing IIDR installation on Clustered Disk.....	22
8.2	Re-using the existing IIDR instance on existing clustered disk .....	23
8.3	Using a new IIDR installation on Local Disk.....	23
8.4	Reusing the existing installation and instance in remote apply topology .....	23
9	Resuming replication after automatic or manual failover to a synchronous replica .....	24
10	Resuming replication after forced failover to an asynchronous replica .....	24

# 1 Deploying IBM IIDR in Microsoft SQL Server AlwaysOn availability group

The following document describes the steps required to be taken in order to deploy IBM IIDR against a database that is participating in a Microsoft SQL Server AlwaysOn Availability Group. The steps are required when the IBM IIDR will serve as a target for data replication. There are 2 deployment options for IBM IIDR as target against databases participating in a Microsoft SQL Server AlwaysOn Availability Group:

- Co-locate IBM IIDR on the same server as the database server. This option is called local apply.
- Deploy IBM IIDR on a different server than the database server. This option is called remote apply.

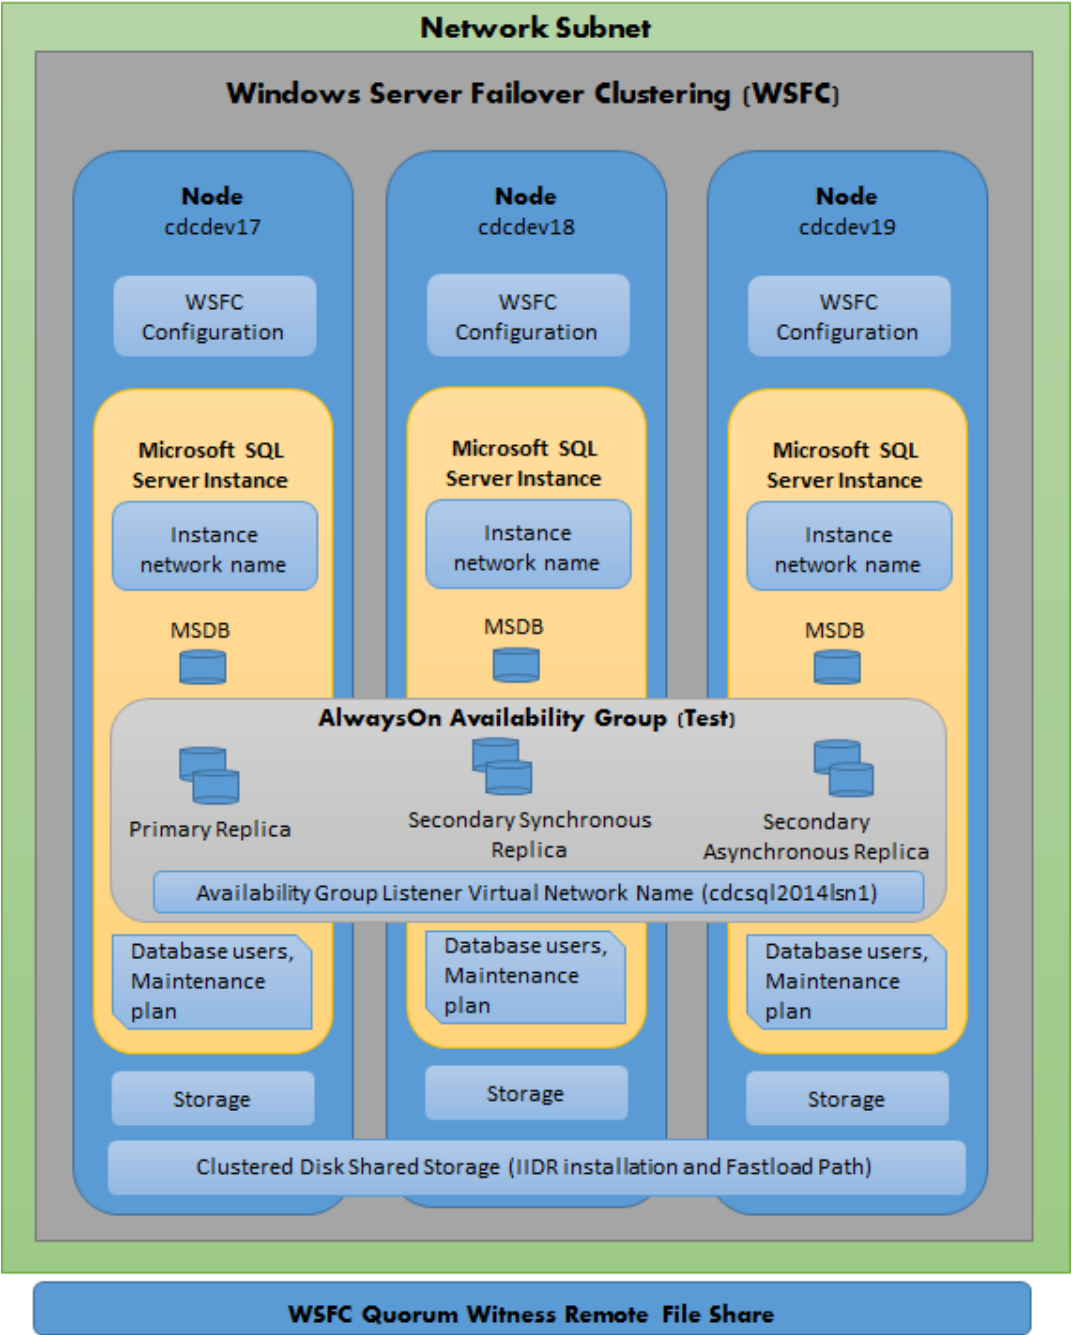
## 2 Local apply

Before proceeding ensure that all prerequisites for installing IBM IIDR are met. Depending on your current circumstances you may choose to deploy IIDR in remote apply or local apply topology. The requirements are different for the different topologies.

### 2.1 Overview

In order to be able to deploy IBM IIDR in a Microsoft SQL Server AlwaysOn local environment one needs to:

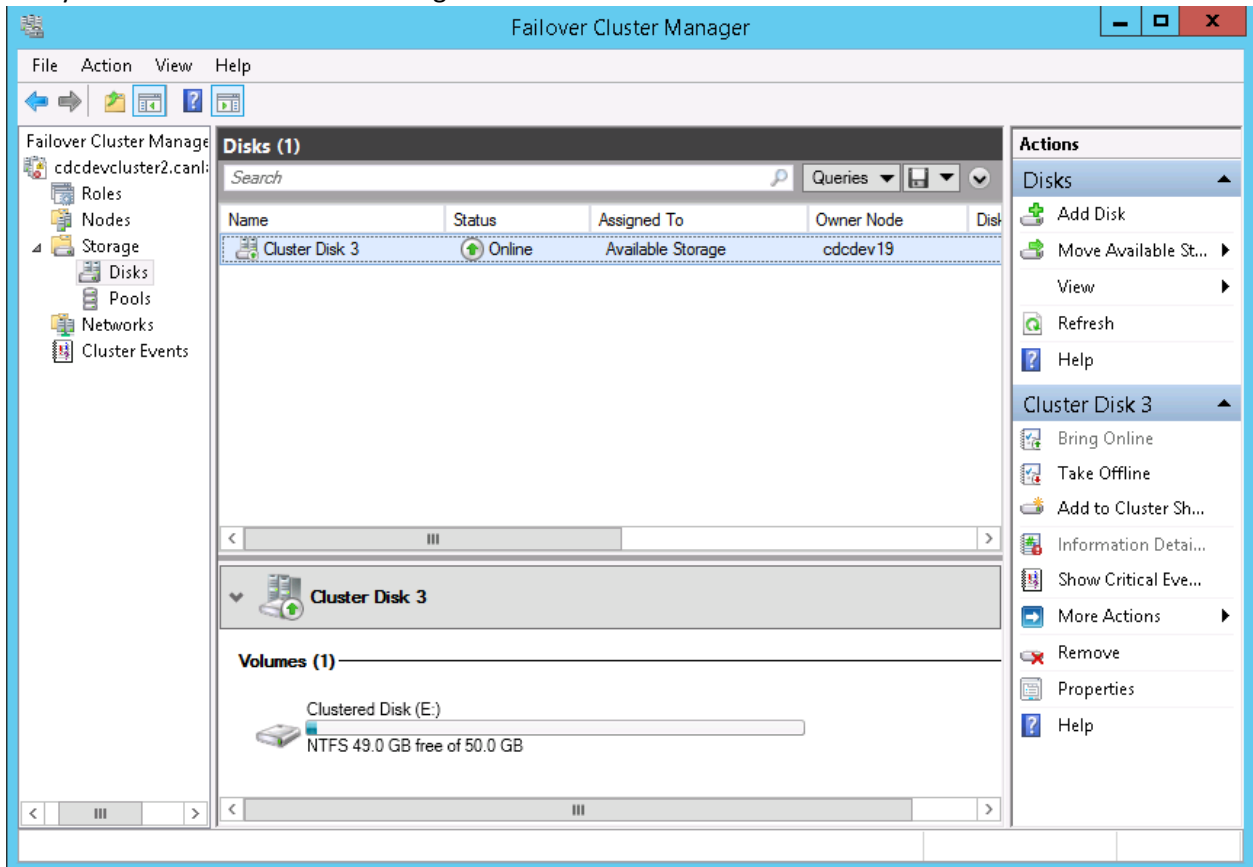
- Add symmetrical shared storage for IBM IIDR deployment reachable from all AlwaysOn nodes.
- Define a windows domain user that can be used to run the IBM IIDR service.
- Create the IBM IIDR user with which it can connect to the database on all nodes where it may fail over.
- Deploy IBM IIDR on the shared symmetrical storage
- Create IIDR service on all nodes to which IIDR may fail during planned or unplanned failover. Any number of nodes that are configured for Synchronous Commit availability mode could be used.
- Add the IBM IIDR service as a dependent resource to the AlwaysOn availability group role that was created in the Windows Failover Cluster.
- Configure the new Datastore in IBM Management Console



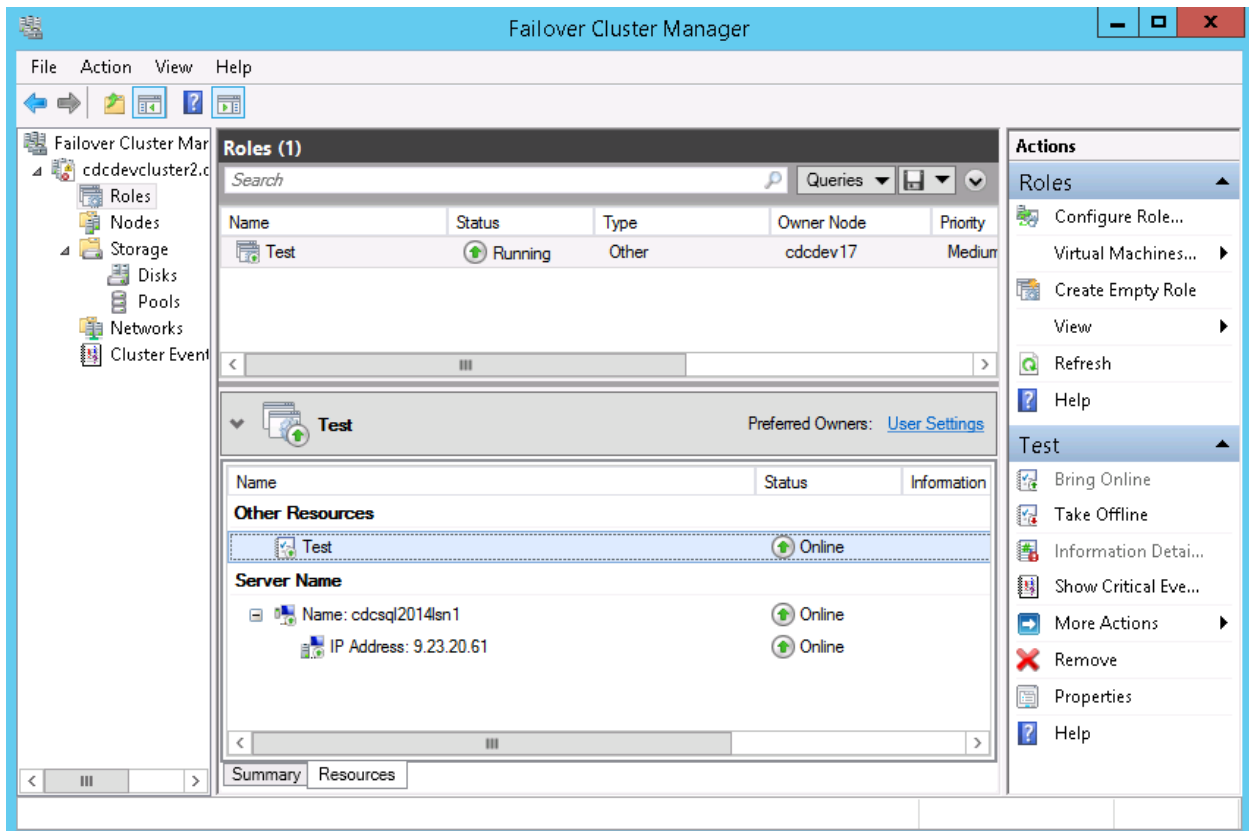
## 2.2 Add symmetrical shared clustered disk reachable from all AlwaysOn nodes

In order to continue with the replication preparation, shared clustered disk that will hold the IBM IIDR binaries and metadata has to be defined as a dependent resource in the AlwaysOn Role that was created in the Windows Failover Cluster for the AlwaysOn deployment. These steps are fairly similar to the steps that are required to be performed when adding shared clustered disk for Microsoft SQL Server Clustering solution: <https://msdn.microsoft.com/en-us/library/ms179530.aspx>. More details on how to add Shared clustered disk can be found at: <http://blogs.msdn.com/b/clustering/archive/2012/06/02/10314262.aspx>


Add your clustered disk to the Storage -> Disks.




Open the failover Cluster Manager and find the Role that was defined by the AlwaysOn availability group configuration.



Right click the AlwaysOn role and select add storage.



Select the disk and click OK.



## 2.3 Define a windows domain user that can be used to run the IBM IIDR service

Given the distributed nature of the AlwaysOn deployment you need to have a windows domain user that is present on all the nodes in the AlwaysOn cluster where IIDR could fail over to.

## 2.4 IIDR Database authentication user

Create the IBM IIDR user with which it can connect to the database on all nodes where it may fail over.

In order to have IBM IIDR fail over smoothly from a primary availability replica to a new primary availability replica, it will need to be able to connect to the database with the same user name and password and similar privileges on any availability replica that it could fail. Now you can define the user with which IBM IIDR will connect to a Microsoft SQL Server instance. This user can be used for data replication or table metadata discovery.

- On each Microsoft SQL Server instance right click Security and select New Login... Choose the user name/password, role as per IBM IIDR user requirements. For more information see:

[http://www-01.ibm.com/support/knowledgecenter/SSTRG7\\_11.3.3/com.ibm.cdcdoc.cdcformssql.doc/concepts/creatingwindowsaccountinsqlserver.dita?lang=en](http://www-01.ibm.com/support/knowledgecenter/SSTRG7_11.3.3/com.ibm.cdcdoc.cdcformssql.doc/concepts/creatingwindowsaccountinsqlserver.dita?lang=en)



## 2.5 Deploy IBM IIDR on the shared clustered disk

Now that we have:

- Added symmetrical shared storage for IBM IIDR metadata.
- Have a windows domain user to run IBM IIDR service under.
- Created database users to be used by IIDR replication.

We can proceed with deploying IBM IIDR product on the AlwaysOn availability group node that holds the current PRIMARY replica for the database. The node that holds the current PRIMARY replica should also be the owner of the shared symmetrical storage where IIDR needs to be deployed. In our example CDCDEV17 is the current primary replica and the owner of the primary replica. The deployment is similar to any deployment of IIDR for Microsoft SQL Server cluster. The difference in deployment is that Microsoft SQL Server does not use the shared storage in this case to store the database and the logs. For more details see:

[https://www.ibm.com/support/knowledgecenter/en/SSTRGZ\\_11.4.0/com.ibm.cdcdoc.cdcforms\\_sql.doc/concepts/configuringclustered.html](https://www.ibm.com/support/knowledgecenter/en/SSTRGZ_11.4.0/com.ibm.cdcdoc.cdcforms_sql.doc/concepts/configuringclustered.html)

Choose the installation directory and the instance directory to be on the symmetrical shared storage.

After deploying the product, one can create an instance. Launch Configuration tool to create the IIDR instance that would replicate data to the database that is part of an AlwaysOn availability group.

The data that needs to be filled out in the configuration tool is not much different from creating an instance in non AlwaysOn environment. A few things to pay attention to:

- For Log on As, choose This account and type in a windows user.
- For host type the AlwaysOn availability group listener name that your database belongs to.
- For Port type the port number for the AlwaysOn availability group.
- Refresh Loader path needs to be a valid path on any node IBM IIDR may fail over to. One option to ensure that is to use the same Shared symmetrical disk used for IIDR installation as the refresh\_loader\_path.

## 2.6 Create IIDR service on secondary synchronous replicas

Do not start the instance after it is created. Similar services that were created on the current node that is responsible for the active PRIMARY replica need to be created on each SECONDARY replica node where IBM IIDR may fail over. Navigate to IBM IIDR installation directory and find the command `dmcreateclusterservice.exe` and invoke the command for a number of replicas where IBM IIDR may fail over.

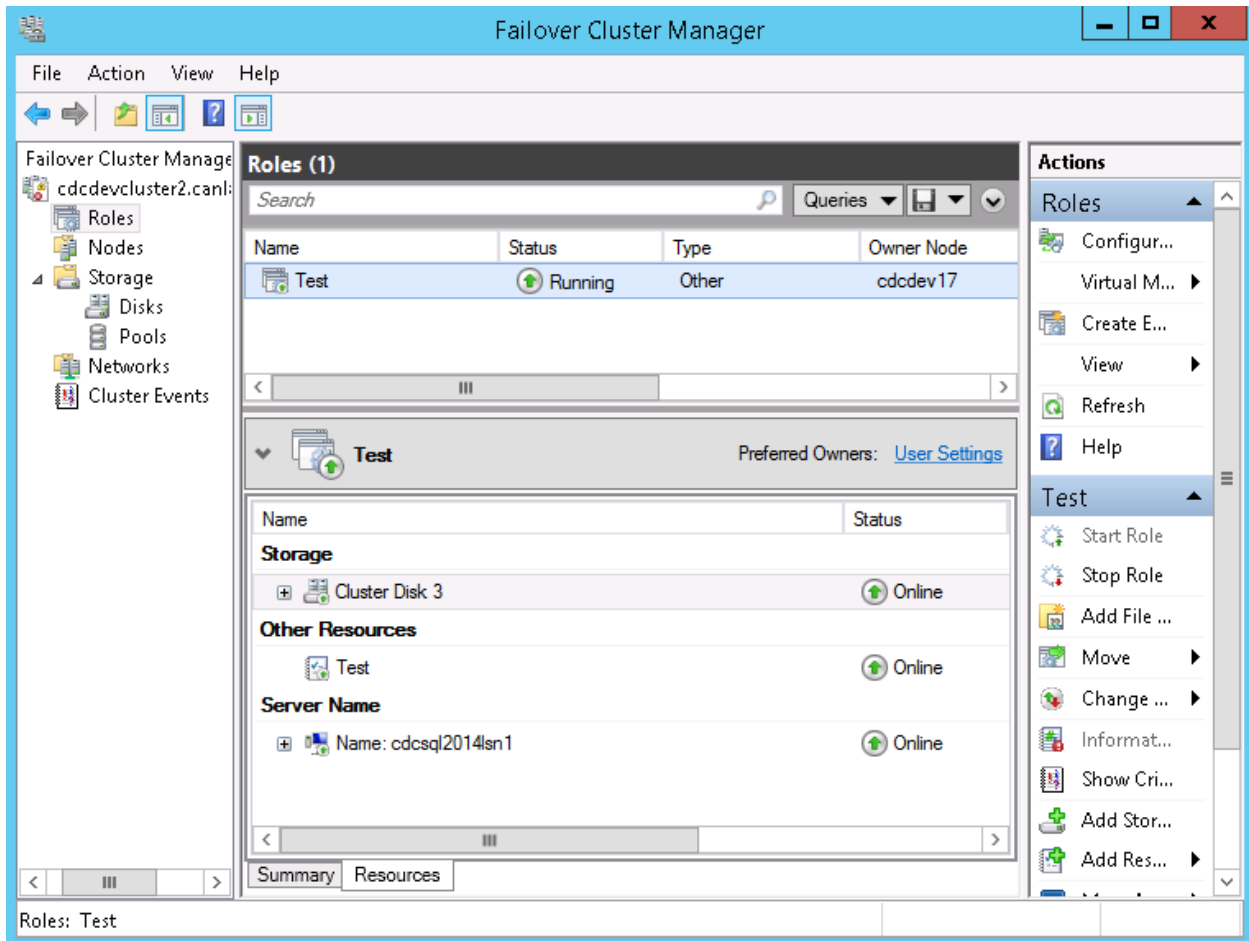
```
dmcreateclusterservice.exe -I <IIDR instance_name> <secondary_replica_node>
```

For the test example that we have we will invoke the command below. IBM IIDR can only fail over to SYNCHRONOUS availability mode replica and the availability group Test is configured to have only CDCDEV18 as a synchronous availability mode replica.

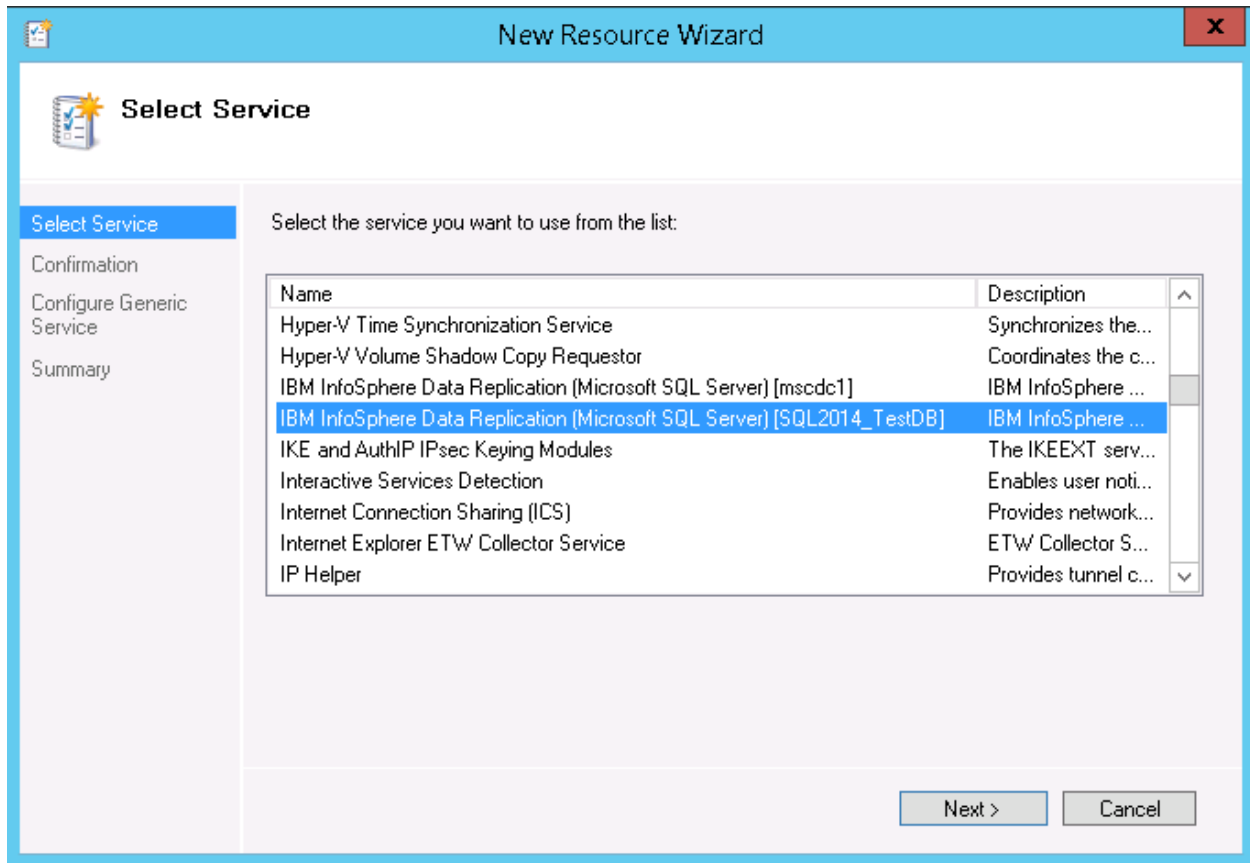
```
dmcreateclusterservice.exe -I SQL2014_TestDb CDCDEV18
```

This command will create a service “IBM InfoSphere Data Replication [SQL2014\_TestDb]” on the `cdcdev18`.

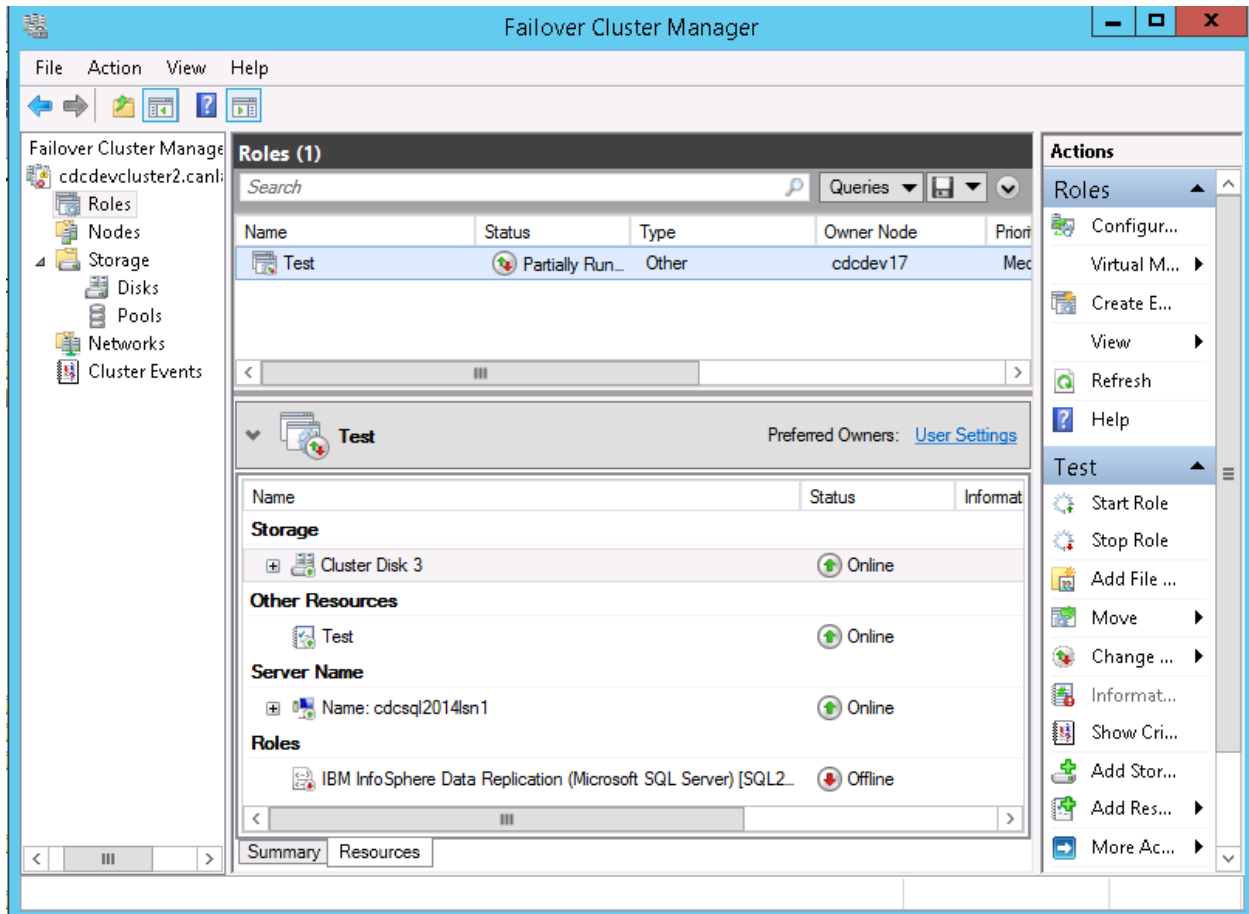
- While being on the PRIMARY node (CDCDEV17) add the IBM IIDR service as a dependent resource to the AlwaysOn availability group role that was created in the Windows Failover Clustering  
The next step includes adding IBM IIDR service as a dependent resource to the AlwaysOn availability group role that was created in Windows Failover Cluster at the time when the availability group was defined.
- Open Windows Failover Cluster manager and navigate to the role created by the availability group.



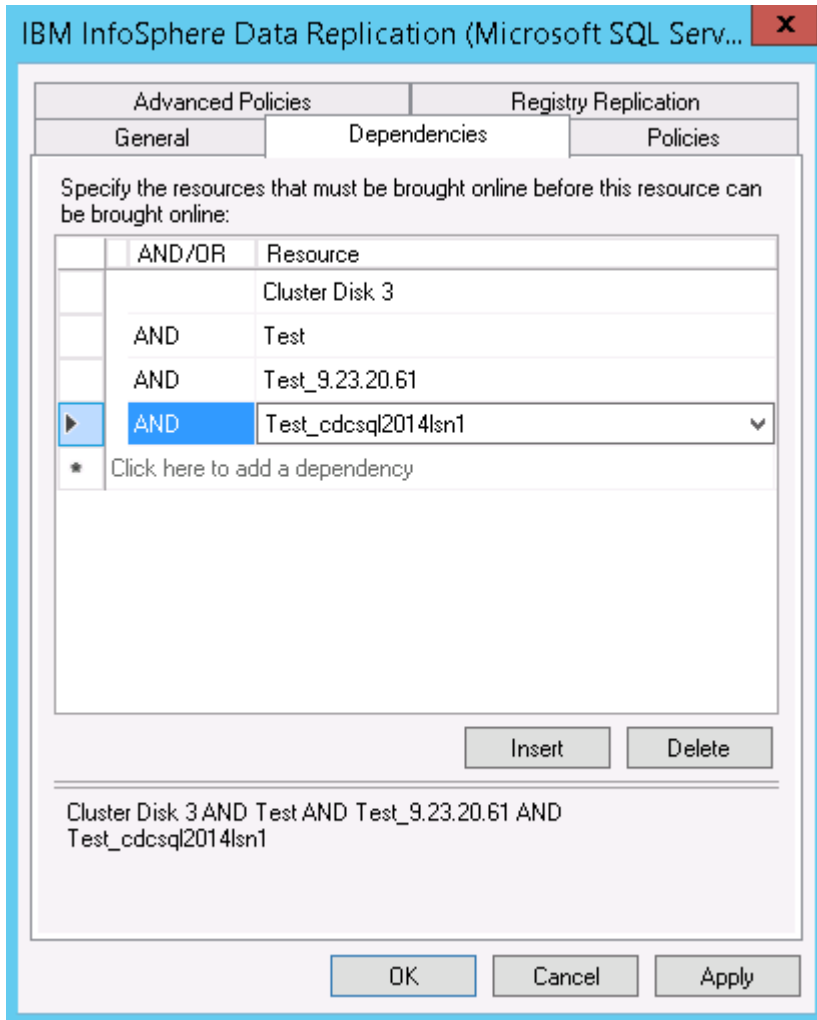
- Right click the Test role and select Add Resource -> Generic Service. Select the service “IBM InfoSphere Data Replication [SQL2014\_TestDb]” and click Next




Click Next in Confirmation dialog. Click Finish in the Summary dialog.



Right click the new resource added to the availability group role and select Properties. Navigate to the Dependencies tab and add the rest of the resources to the list of dependencies. This will prevent IBM IIDR resource from being started prior to having the rest of the resources initialized and started after a failover.



Navigate to the Policies tab and choose the policy that you would want to be applied when the service fails or is manually stopped. In most cases customers may not want to fail over the entire availability group in case IBM IIDR fails or is manually stopped. If this is the case select “If Resource Fails, do no restart”. For more information about resource properties consult Microsoft documentation. Click OK.



Right click the service and select “Bring Online”

### 3 Remote apply

Before proceeding ensure that all prerequisites for installing IBM IIDR are met. Depending on your current circumstances you may choose to deploy IIDR in remote apply topology.

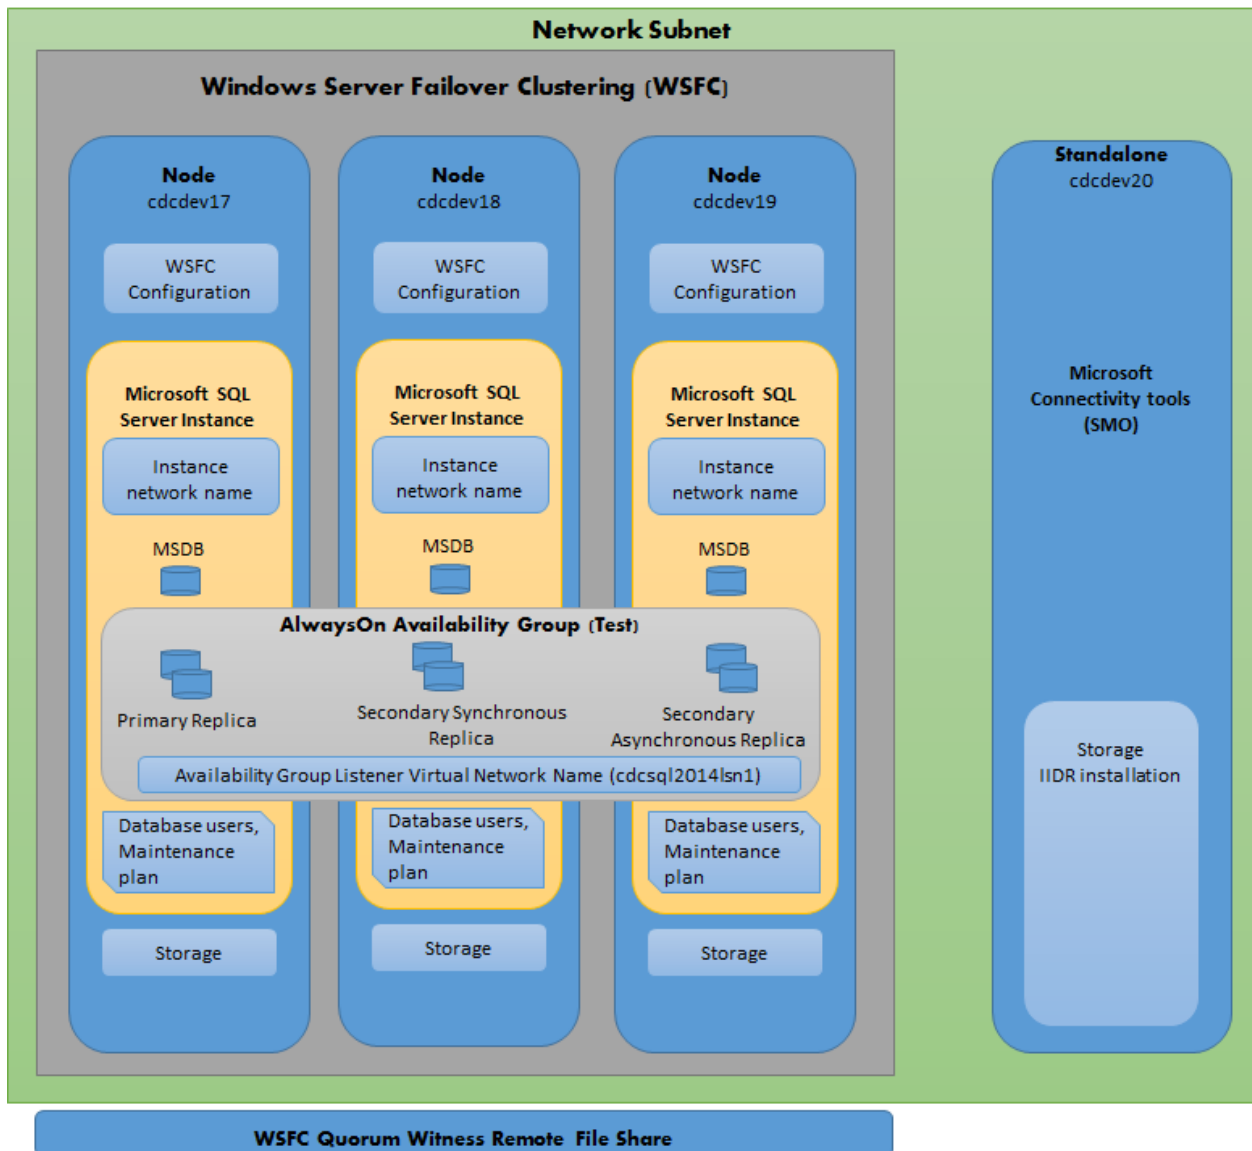
#### 3.1 Overview

In order to be able to deploy IBM IIDR in a Microsoft SQL Server AlwaysOn remote environment one needs to:

- Define a windows domain user that can be used to run the IBM IIDR service.



- Install the Microsoft connectivity tools (SMO) on the server where you plan to install IIDR. For more information check the IIDR documentation.
- Install Microsoft SQL Server connectivity tools on the server where you install InfoSphere CDC. For more information check the product documentation pre-requisites for IBM IIDR for Microsoft SQL Server remote apply. [http://www-01.ibm.com/support/knowledgecenter/SSTRGZ\\_10.2.1/com.ibm.cdcdoc.cdcformssql.doc/concepts/configuringaremotetarget.html?lang=en](http://www-01.ibm.com/support/knowledgecenter/SSTRGZ_10.2.1/com.ibm.cdcdoc.cdcformssql.doc/concepts/configuringaremotetarget.html?lang=en)
- Create the IBM IIDR user with which it can connect to the database on all Microsoft SQL Server instances where the availability group may fail over.



### 3.2 Deploy IBM IIDR in remote apply topology

Now that we have:

- Have a windows domain user to run IBM IIDR service under.
- Created database users to be used by IIDR replication.
- Installed the Microsoft SQL Server connectivity tools (SMO)
- We can proceed to install IIDR in remote apply topology. The installation for deployment in AlwaysOn remote apply topology is fairly similar to the remote apply against a database not residing in an AlwaysOn availability group. In order to deploy the product as an Early Release product evaluation that would be AlwaysOn aware one needs to install the product with the AlwaysOn enabling knob using the command line:  
setup-iidr-11.3.3.{x}-{y}-x86-sql.exe -DPRERELEASE=alwayson
- After deploying the product, one can create an instance. Launch Configuration tool to create the IIDR instance that would replicate data to the database that is part of an AlwaysOn availability group.

IBM InfoSphere Data Replication New Instance

Instance Configuration AlwaysOn Communications Protocol

Instance

Name: SQL2014\_TestDB

Server Port: 10501

Staging Store Disk Quota (GB): 100

Maximum Memory Allowed (MB): 1024

Windows Service

Log on as:

Local System Account

This account: MYDOMAIN\myuser

Password: .....

Confirm Password: .....

Database

Host: cdcsql2014lsn1

Name: TestDb Advanced...

Port: 5025

Metadata Schema: metadata ...

Database Authentication

SQL Authentication

User: sql\_server\_authentication\_user

Password: .....

Windows Authentication

Refresh Loader

Refresh Loader Path: {SharedDrive}:\refresh\_loader\_path Browse...

OK Cancel

The data that needs to be filled out in the configuration tool is not much different from creating an instance in non AlwaysOn environment. A few things to pay attention to:

- For Log on As, choose This account and type in a windows user.
- For host type the AlwaysOn availability group listener name that your database belongs to.
- For Port type the port number for the AlwaysOn availability group.
- Refresh Loader path needs to be a writable by IIDR service user and readable by Microsoft SQL Server service user. The path should also be symmetrically accessible on both IIDR server and any node to which the AlwaysOn availability group could fail over to.

## 4 Configure the new Datastore in Management Console

When creating a new Datastore in Management Console use the AlwaysOn availability group listener name as the host name for the new Datastore. In our example use `cdcsql2014lsn1` as the host name for the new Datastore. The rest of the configuration is similar to defining a new Datastore. For more information look up IBM IIDR documentation on configuring a new Datastore.

## 5 Test the deployment

In order to test the deployment one would need to:

- Create a test source table in the source database
- Create a test target table in the AlwaysOn target database
- Create a test subscription in Management Console and create a table mapping between the newly created source test table and a target table.
- Insert a couple of rows in the source test table.
- Start mirroring.
- Check that the table has been successfully refreshed and is mirroring.
- Insert a few more rows. Test that replication has been successful.
- Start Windows Failover Cluster Manager and navigate to the AlwaysOn availability group role.
- Right click on the role and select Move -> Select Node. Select the Synchronous replica node (CDCDEV18)
- All resources will fail over and IBM IIDR will be started.
- The subscriptions will not be automatically started. The user will have to manually start the subscriptions or use persistent subscriptions that permits restarting subscriptions after a recoverable error. Check your source datastore documentation on the level of support for persistent subscriptions
- Start the subscription. Insert a few more rows. Check that data has been replicated to the target.

## 6 Persistent subscriptions

Persistent subscriptions permit subscription restart after a recoverable error. Persistent subscriptions is a feature of a source datastore. Check the source datastore documentation to evaluate the support for persistent subscriptions in the source datastore. If the persistent subscription feature is available for your source datastore, it could be used in AlwaysOn deployments in order to automatically restart the subscriptions after a failover. In order to enable persistent subscriptions open IBM Management Console, connect to the source datastore, right click on the subscription, select properties, then Advanced Settings and tick the “Mark Subscription as persistent”.

## 7 Considerations when transitioning the database to an AlwaysOn availability group

As AlwaysOn support requires an IIDR deployed on a shared clustered disk and this is not likely to be a deployment for non-AlwaysOn deployment, a simple upgrade will not allow a deployment to be converted to support AlwaysOn. A new IIDR deployment is needed in order to support AlwaysOn. One needs to follow the instructions in this document in order to deploy a new IIDR installation that would support AlwaysOn.

### 7.1 Using a new IIDR installation and instance

If you currently have IIDR deployed and your database is not AlwaysOn, the following steps need to be executed in order to add the database to AlwaysOn availability group and to deploy IIDR that supports such a database deployment.

- Establish a quiesce point for the source database. This means that the database cannot be accessible to external transactions either DML or DDL. If a quiesce point cannot be established a bookmark transfer will be needed
- Replicate data to the head of log.
- Upgrade your existing IIDR instance to the latest version that supports AlwaysOn. Start the instance. This will upgrade the IIDR metadata to the latest product version.
- Upgrade Management Console and Access Server to the latest version.
- Add the database to an AlwaysOn availability group.
- Deploy a new IIDR installation on the shared clustered disk and follow all the steps required for an AlwaysOn IIDR deployment.
- Start both instances.
- Use Management Console to either export/import subscriptions or promote subscriptions from old non AlwaysOn deployment to AlwaysOn deployment.
- If a quiesce point has been established at the beginning and all data has been replicated to the head of log prior to upgrade, you can Mark Table Capture Point for Mirroring for all tables in all subscriptions in the AlwaysOn instance. This action will tell IIDR to start reading from the current head of log. If a quiesce point could not be established at the beginning, then you can transfer the bookmark from the old subscriptions to the new subscriptions.
- A good rule of thumb is Test and document the specific steps needed for your solution in a test environment prior to making changes to production environment.

### 7.2 Reusing the existing installation and instance in remote apply topology

- Stop replication from all sources
- Stop the target instance

- Add the database to the availability group.
- Start the configuration tool and change the host and port from the Microsoft SQL Server host and port to the availability host and port.
- Refresh Loader path needs to be a valid path on any node IBM IIDR may fail over to. One option to ensure that is to use the same Shared symmetrical disk used for IIDR installation as the refresh\_loader\_path.
- Start IIDR. Start replication.
- A good rule of thumb is Test and document the specific steps needed for your solution in a test environment prior to making changes to production environment.

## 8 Considerations when transitioning the database out of an AlwaysOn availability group

One may decide to continue to use IIDR replication and no longer have their database as part of an availability group. In order to achieve this one needs to deploy a new IIDR instance configured against the non AlwaysOn database. Here are the steps that need to be executed when transitioning from a database that is part of an AlwaysOn availability group to the same database that is not part of an availability group.

### 8.1 Using existing IIDR installation on Clustered Disk

The following steps need to be executed when one wants to transition the database out of an availability group and change IIDR to use new binaries residing on clustered symmetrical clustered disk.

- Establish a quiesce point for the source database. This means that the database cannot be accessible to external transactions either DML or DDL. If a quiesce point cannot be established a bookmark transfer will be needed
- Replicate data to the head of log.
- Remove the target database from the AlwaysOn availability group.
- Create a new IIDR instance against the database that is not part of an availability group.
- Start both instances.
- Use Management Console to either export/import subscriptions or promote subscriptions from old AlwaysOn instance to non AlwaysOn instance.
- Remove the IIDR service of the AlwaysOn instance from the availability group.
- If a quiesce point has been established at the beginning and all data has been replicated to the head of log prior to upgrade, you can Mark Table Capture Point for Mirroring for all tables in all subscriptions in the AlwaysOn instance. This action will tell IIDR to start reading from the current head of log. If a quiesce point could not be established at the beginning, then you can transfer the bookmark from the old subscriptions to the new subscriptions.

- A good rule of thumb is Test and document the specific steps needed for your solution in a test environment prior to making changes to production environment.

## 8.2 Re-using the existing IIDR instance on existing clustered disk

- The steps for reusing the existing IIDR instance are similar to the steps for reusing the existing installation [remote apply](#) deployment topology.

## 8.3 Using a new IIDR installation on Local Disk

The following steps need to be executed when one wants to transition the database out of an availability group and move the CDC binaries from the clustered shared disk to local disk.

- Establish a quiesce point for the source database. This means that the database cannot be accessible to external transactions either DML or DDL. If a quiesce point cannot be established a bookmark transfer will be needed
- Replicate data to the head of log.
- Upgrade your existing AlwaysOn IIDR instance to the latest version of the product that you want to deploy on local drive. Start the AlwaysOn instance. This will upgrade the IIDR metadata to the latest product version.
- Upgrade Management Console and Access Server to the latest version.
- Remove the database from an AlwaysOn availability group.
- Deploy a new IIDR installation on the local disk.
- Start both instances.
- Use Management Console to either export/import subscriptions or promote subscriptions from old AlwaysOn deployment to the new non AlwaysOn deployment.
- If a quiesce point has been established at the beginning and all data has been replicated to the head of log prior to upgrade, you can Mark Table Capture Point for Mirroring for all tables in all subscriptions in the AlwaysOn instance. This action will tell IIDR to start reading from the current head of log. If a quiesce point could not be established at the beginning, then you can transfer the bookmark from the old subscriptions to the new subscriptions.
- A good rule of thumb is Test and document the specific steps needed for your solution in a test environment prior to making changes to production environment.

## 8.4 Reusing the existing installation and instance in remote apply topology

- Stop replication from all sources
- Stop the target instance
- Start the configuration tool and change the host and port from the availability group host and port to Microsoft SQL Server instance that will hold the database after it is removed from the availability group. Make sure that the rest of the configuration continues to make sense for the

standalone database that will be made available after the database is removed from the AlwaysOn availability group. Save.

- Remove the database from the availability group. Make sure that the only remaining instance of the database is present on the Microsoft SQL Server that was configured in IIDR configuration tool.
- Start IIDR. Start replication.
- A good rule of thumb is Test and document the specific steps needed for your solution in a test environment prior to making changes to production environment.

## 9 Resuming replication after automatic or manual failover to a synchronous replica

As you probably have already tested, after a failover to a new PRIMARY replica, IIDR fails together with the availability group to the new node in local apply deployment topology and loses the connection to the database in the remote apply deployment topology. Assuming the subscriptions have been marked persistent on the source datastore, the replication will also automatically restart and continue from where it left over on the previous PRIMARY replica.

## 10 Resuming replication after forced failover to an asynchronous replica

IBM IIDR only supports replication to a synchronous replica. If there was a need to fail over the database to an asynchronous replica, then IIDR is not expected to successfully replicate data. IIDR does not support failing over to an asynchronous replica. In addition, if you failover back to the previous PRIMARY replica, IIDR will not be able to continue mirroring, because IIDR does not support failover from an ASYNCHRONOUS replica.

In order to recover from such a situation, you will need to either define a new set of synchronous replicas or fail over to the original PRIMARY replicas, review that Microsoft Replication continues to be properly configured for the environment and initiate Refresh/Mirror for all tables.