

Configuring InfoSphere Remote Source for Db2/z for a resilient environment

Table of Contents

- Overview..... 1**
 - Installation and configuration steps..... 3
 - Failover Steps 3
- CDC Instance Configuration 3**
 - Automatic configuration option 3
 - Manual configuration option 4
 - Other considerations: 4
- Failover Scenarios..... 5**
 - LPAR failure 5
 - DB2 member from the data sharing group failure 5
 - VMWare node failure 5
 - InfoSphere CDC instance failure 6
- Other types of HA configurations..... 6**

Overview

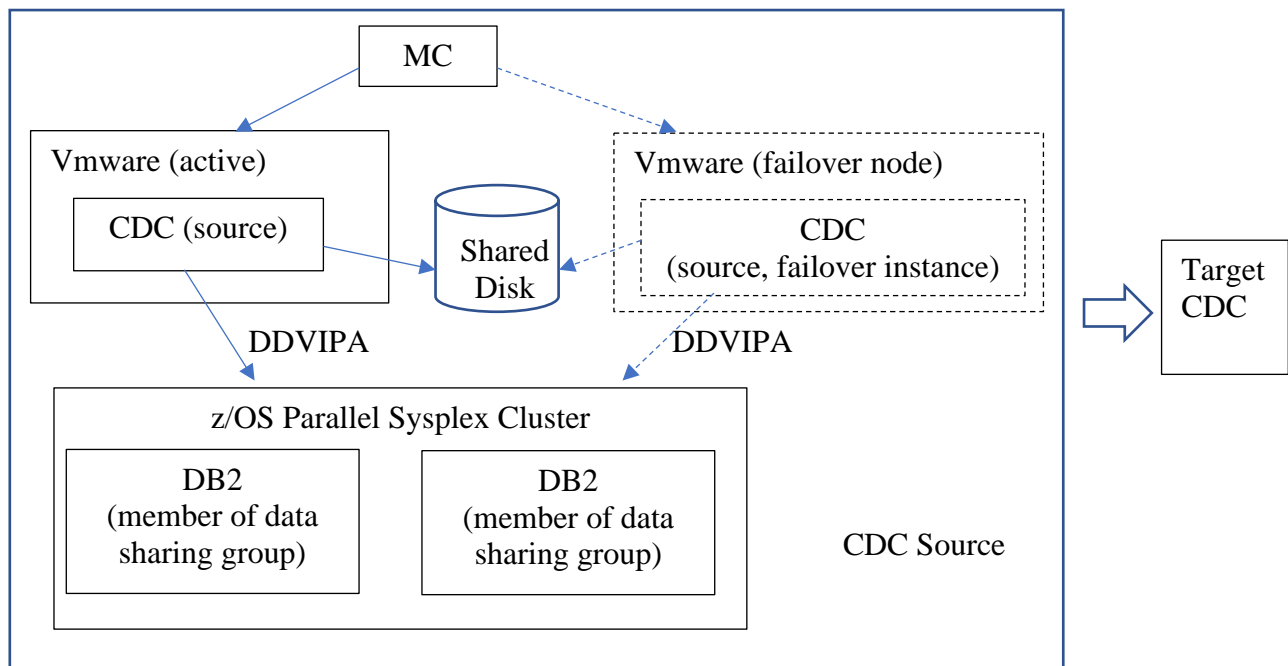
CDC Remote Source for Db2 z/OS is a client application that interacts with Db2 on z/OS through standard JDBC interfaces, using a product-supplied stored procedure to obtain log data. This document describes High Availability (HA) considerations of CDC remote capture sourcing Db2 in a data sharing group on a Parallel Sysplex cluster. The objective is for replication to quickly resume normal operations should any replication component(s) fail.

The Parallel Sysplex architecture provides capabilities for disaster recovery from LPAR failures. The Db2 Data Sharing Group is typically implemented for high availability for Db2. This means that applications connecting to Db2 through a shared virtual IP/port (typically DDVIPA), for example, can immediately reconnect successfully after connection loss due to an LPAR or a Db2 subsystem failure.

CDC remote capture must be implemented in a Linux, Unix, or Windows environment. A typical HA deployment configuration (for all CDC luw agents) is to deploy a CDC remote capture agent on two separate virtual or real machines, with a shared disk where the instance data and metadata is stored. One of these two machines will be the active machine, while the other is the passive, or standby, machine. This is a common Cold/Standby HA configuration for all LUW source and target agents. The active CDC instance connects to the mainframe cluster through a shared IP address (typically a DDVIPA).

This document describes the typical implementation using VMWare virtual machines and automated failover capabilities from VMWare vMotion. Note: VMWare is not required. Customers can employ their own external automation to start the standby instance after failure detection. Other clustering systems can also be used for this purpose, such as Kubernetes (for container versions of CDC).

The architecture is shown below. The components in solid boxes are the ones in active use; the components in dashed-line boxes become active should the current components fail.



When active instance fails and replication is switched to the standby, two main points should be considered:

1. How does Management Console connect to the standby CDC instance?
2. How does the standby CDC instance find the configuration and operational metadata?

Management Console should use the virtual IP defined for the VMWare cluster to connect to the active CDC instead of the physical IP of any VMWare node.

For CDC metadata to be accessible to the failover node, ensure that the failover node is mounted to the same \$CDC_HOME path, and that the CDC user/group is available on the passive node with the same permissions.

Installation and configuration steps

1. On the active VMWare node, install CDC and configure CDC instances/subscriptions.
2. From Management Console, connect to the active VMWare with the virtual IP defined for the VMWare cluster.

Failover Steps

If failover is planned, do the following on the active node. Skip the steps if failover is not planned

1. Stop replication for all subscriptions (run **dmendreplication**).
2. Stop InfoSphere CDC instance (run **dmshutdown**).

Recovery steps on the failover node, after the shared volume becomes available,

1. Start the CDC instance (run **dmsts64**).
2. Clear the staging store for the instance (run **dmclearstagingstore**).
3. Restart subscriptions (run **dmstartmirror**)

CDC Instance Configuration

Automatic configuration option

https://www.ibm.com/support/knowledgecenter/SSTRGZ_11.4.0/com.ibm.cdcdoc.cdcd2zremote.doc/concepts/configoptions_automatic.html

When using CDC's automatic configuration tool (**dmconfigurets**) to create a new instance or upgrade an existing instance, the Remote Source for Db2/z CDC agent requires the following connections:

1. SSH connection to make changes to the z/OS environment for creating external stored procedure, such as copying the object modules, link-edit the load module, and refreshing the WLM environment. This connection is used at installation time and at upgrade time.
2. JDBC connection to execute Db2 sql statements such as retrieving schema information and creating the external stored procedure. It is also used at runtime for reading from IFI 306 via the stored procedure. The IP used for JDBC connection in a high availability environment should be the DDVIPA virtual IP address.

(Note: Both #1 and #2 use different sessions but can share the same hostname/IP address. Normally, this requires the use of a shared (eg. DDVIPA) virtual IP address.)

After the one-time setup, the SSH connection is no longer required. CDC replicates data using only the JDBC connection.

The automatic configuration tool requires an input for database hostname/IP address. It is recommended to use the DDVIPA address defined for both DDF and SSH connections.

However, some customers may have enabled only one member from the Db2 Data Sharing Group to allow SSH connection through the member's physical IP address. In this case, use the member's physical address for the database hostname/IP address field in the CDC instance creation prompt, and configuration will succeed. Then change to the DDVIPA IP before starting this instance to enable HA capability.

When upgrading the CDC instance, use the member's physical address to configure the instance. The instance must be started for the upgrade to take place. After the upgrade is complete, the instance may fail due to connection error for JDBC. Then change back to the DDVIPA IP and restart the instance.

Manual configuration option

https://www.ibm.com/support/knowledgecenter/SSTRGZ_11.4.0/com.ibm.cdcdoc.cdcd2zremote.doc/concepts/configoptions_manual.html

In the case where no SSH connection is enabled, or that the DDVIPA IP is separate from the SSH IP address, use the manual configuration steps outlined in the above documentation.

Start by attempting to create an instance using `dmconfigurets` with the DDVIPA IP address, and when SSH steps fail CDC generates an output file called `deploy.txt` with instructions to setup manually. The precise manual steps would need to be modified for the specific environments since the manual steps refer to SSH. After the manual steps are completed, the external stored procedure should be created successfully for the database/schema combination. Rerun `dmconfigurets` tool, and if the stored procedure exists and gives back good information when CDC calls it, then the configuration will continue successfully.

If automatic upgrade fails with SSH connection, CDC generates a new `deploy.txt` file so that you can upgrade manually.

Other considerations:

- When defining the instance manually, each instance should have its own schema and WLM environment.
- Mark the subscriptions "persistent" so that the subscriptions will automatically restart if a transient problem occurs to the subscription.

- It is recommended to use shared scrape so that all subscriptions share the same JDBC session for efficient use of CPU. This can be configured by setting the CDC parameter **use_staging_store** to **true** (default is false for remote capture, true for all the other CDC engines). This means that if the JDBC connection fails connecting to one Db2 member, then all subscriptions fail. However, in a Data Sharing Group environment, the system will automatically switch to a surviving member and the JDBC connection will resume. If the subscriptions are configured for persistency, they will automatically resume as well.

Failover Scenarios

The four main resiliency scenarios are:

1. LPAR failure
2. Db2 Data Sharing Group member failure
3. VMWare node failure
4. InfoSphere CDC instance failure

LPAR failure

For the failing instance, all in-flight transactions and sessions, including any that CDC agent have with that instance, are rolled-back and/or lost. Applications will retry connections and when they do, they will be serviced by a surviving Db2 instance on another LPAR in the SYSPLEX.

Replication is, in that sense, just another database application. If the CDC session is lost, the subscription is stopped. Once restarted, CDC will connect to the DBMS and to the stored procedure on another instance via the DDVIPA. In a typical Db2 data sharing environment, all tables and stored procedures are shared across multiple db2 instances in the plex. Replication will continue after subscription restart.

DB2 member from the data sharing group failure

If a member from the data sharing group is lost where your JDBC session is connected to, subscriptions will go down. If subscription persistency is enabled and CDC is connected through DDVIPA, the affected subscriptions will automatically reconnect through a different JDBC connection. (Note: session loss can happen for other reasons as well, such as noise bursts and network device failures)

VMWare node failure

If the VMWare node is configured for HA in a clustering environment (using VMotion, for example), the cluster will automatically switch to a different VMWare should the active one fails. To achieve automatic persistency, the failover steps outlined in the [recovery steps on the failover node](#) should be part of the startup script for the VMWare node.

InfoSphere CDC instance failure

This is the case where CDC instance fails but the VMWare node is still active. CDC does not provide built-in functionality to mark an instance “persistent” and automatically restart the instance (this functionality is only available for subscriptions). To achieve automatic persistency, an external script can be implemented to detect if CDC instance is alive and restart the instance. This can solve transient problems such as unstable network to the mainframe host. However, if the CDC instance is down due to a non-recoverable problem, which happens rarely, please open an IBM support ticket.

There are many ways to detect whether or not CDC is alive. For example

```
[cdc@mavin1 bin]$ cat alive.sh
#!/bin/bash
echo "check if instance 'cdcz' is alive"
./dmconfigurets << EOF > /tmp/output
1
6
EOF
grep cdcz /tmp/output
rm /tmp/output
[cdc@mavin1 bin]$ ./alive.sh
check if instance 'cdcz' is alive
cdcz      11801      STLEC1      SYSADM      running
```

Other types of HA configurations

The HA implementation described above is the topology for CDC on a shared volume. There are other types of topologies such as CDC on separate nodes with a shared database, and CDC on separate servers with separate databases. The implementation and recovery steps are different in these scenarios. You can find more details from IBM Redbook,

["Smarter Business Dynamic Information with IBM InfoSphere Data Replication CDC"](#), section 7.5.