

## Vault Project - Plant Database Replication

This document describes how to replicate the plant database for a vault project between WAN connected locations. By replicating both the vault and the plant database to a remote office AutoCAD Plant 3D 2016 and AutoCAD P&ID desktops can maintain a fast LAN-only server connection.

This document also describes where to locate the LOG files that can be used to monitor replication, including how to identify synchronization issues (for example: a collision from two users that check-in using the same tag) between locations.

After the plant database is replicated, run **plantvaultdatabase** at the Command: prompt in AutoCAD Plant 3D or AutoCAD P&ID. The command specifies which plant database to use when you open a vault project from a particular vault server. See product help for more information: [PLANTVAULTDATABASE](#).

## Contents

|  |    |
|--|----|
| Vault Project - Plant Database Replication .....   | 1  |
| Software Requirements: .....                       | 1  |
| New in 2016 .....                                  | 2  |
| Vault Projects and the Plant Master Database ..... | 2  |
| Plant Database Replication .....                   | 4  |
| Synchronization Service .....                      | 7  |
| SQL Server Jobs – Replication Limitation .....     | 11 |
| Replication Topology and Conflicts.....            | 12 |
| Disabling Replication in SQL Server .....          | 13 |
| Maintenance and Manual Audit .....                 | 14 |
| Project Setup - PLANTVAULTDATABASE .....           | 15 |

## Software Requirements:

- Autodesk Vault 2016 replicated vault servers
- Microsoft SQL Server 2008 R2 or later
- SQL Server Agent must be running
- PlantDatabaseReplication.exe (from AutoCAD Plant 3D 2016)

## New in 2016

AutoCAD Plant 3D 2014 and AutoCAD P&ID 2014 and later support vault projects, but did not support replication until 2016 (available Spring 2015).

With AutoCAD Plant 3D 2015 Extension 1 (Summer 2014) vault server replication is supported, but not plant database replication (this meant that the remote desktop needed a WAN connection to the plant database).

AutoCAD Plant 3D 2016 and AutoCAD P&ID 2016 support both vault and plant database replication.


Performance for both vault and non-vault projects that connect to SQL Server over a WAN has been improved with 2016. For vault projects, you will experience significantly faster check-in and check-out times. For non-vault projects, the plant database is updated faster when you save your drawings.

**Note:** You may discover that remote office performance has improved such that you do not need to replicate the plant database compared to AutoCAD Plant 3D 2015 Extension 1.

## Vault Projects and the Plant Master Database

Before describing how to set up plant database replication, it is important to understand that replicating the plant database is optional. Even if you are replicating vault servers, you are not required to replicate the plant database.

With all plant vault projects, AutoCAD Plant 3D and AutoCAD P&ID create a local workspace that you work in on your desktop computer. This means that as you work in your 3D model or P&ID drawings performance is not limited by a WAN connection to the servers.




**No Replication - Neither the vault nor the plant database is replicated**

Shown above, the remote workstations are connected over VPN. You must be connected to the servers to check-in and check-out files, but otherwise you work in your local workspace.

When you create a vault project, you specify the SQL Server to create the plant database on. This is the Default Plant Database server.


You can *replicate* the vault, but decide to *not replicate* the plant database for a vault project.



**Vault-only replication - The vault is replicated but the plant database is not replicated**

In this scenario, the remote workstation has a LAN connection to a replicated vault but a WAN connection to the default plant database. Files are downloaded from and uploaded to the vault over a LAN connection. If performance is acceptable, it is not necessary to replicate the plant database.


However, if performance is slow, or a WAN connection from the workstation to the plant database is not available you can replicate the plant database (which is the subject of this document).



**Full replication – Both the vault and plant database are replicated**


If you replicate both the vault and plant database, remote office workstations do not need to connect to any server using a WAN connection. Only the replicated servers maintain a WAN connection to each other.

## Plant Database Replication



PlantDatabaseReplication.exe is run on each of the client plant database servers. The service keeps each plant database synchronized. It is recommended to run the service on each client plant database server where it maintains a WAN connection to the default plant database.

Before running as a synchronization service, PlantDatabaseReplication.exe is first used in manage mode to copy the plant database.



**PlantDatabaseReplication.exe** and support files can be copied from an AutoCAD Plant 3D desktop installation to the client server. The files listed below are found in the program folder in the default

locations C:\Program Files\Autodesk\AutoCAD 2016 and C:\Program Files\Autodesk\AutoCAD 2016\PLNT3D or C:\Program Files\Autodesk2016\AutoCAD 2016\PNID.

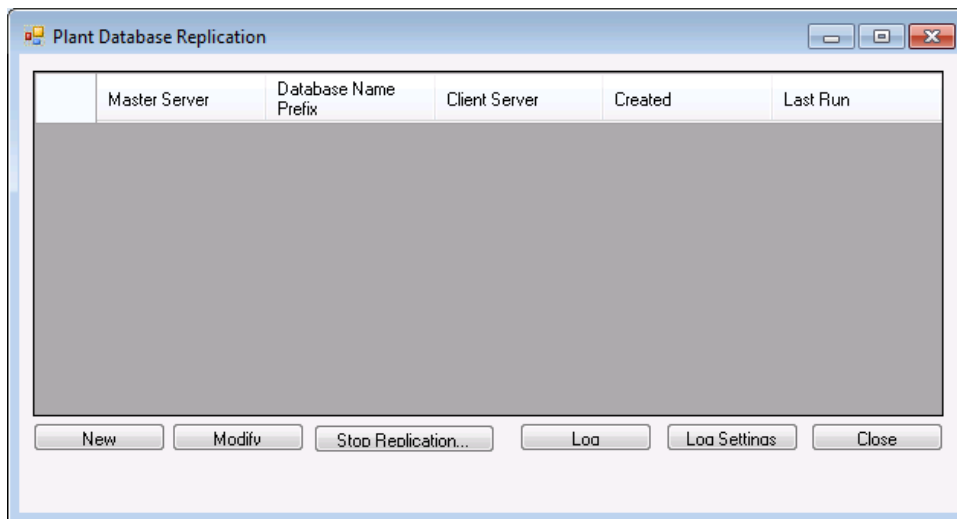
1. PlantDatabaseReplication.exe
2. PlantDatabaseReplication.exe.config
3. PnPCommonMgd.dll
4. PnPDataObjects.dll
5. PnPReplicator.dll
6. PnPSQLiteEngine.dll
7. System.Data.SQLite.DLL
8. log4net.DLL

The following text can be used to create a BAT file which will copy from the default program files location to C:\Replication:

```
robocopy "C:\Program Files\Autodesk\AutoCAD 2016" "C:\Replication" PnPCommonMgd.dll
PnPDataObjects.dll PnPSQLiteEngine.dll System.Data.SQLite.DLL
robocopy "C:\Program Files\Autodesk\AutoCAD 2016\PLNT3D" "C:\Replication" log4net.dll
PlantDatabaseReplication.exe PlantDatabaseReplication.exe.config PnPReplicator.dll
```

[Visual C++ Redistributable for Visual Studio 2012 Update 4](#) is required ().

At the command prompt, run: **PlantDatabaseReplication.exe MANAGE**



In manage mode, you can create a new replicated database.

Click New to display the Create and Replicate Plant Database dialog box.

**Create and Replicate Plant Database**

**Master Plant Database**

SQL Server: 144.111.19.200

Database name prefix: integration\_testing\_1

☒ Windows Authenticaion

User name:

Password:

Test Connection

**Client Plant Database**

SQL Server: PLANT-WAN

Database name prefix: integration\_testing\_1

☒ Windows Authenticaion

User name:

Password:

Test Connection and Name

Check for updates: 5 minutes

Create Client Database Cancel

Specify the default plant database as the master plant database. The master database is usually the default database. Click Test Connection to test the connection to the master database.

In the database name prefix, select the plant database.

Next, specify the client database that you want to replicate or enter *localhost*. Click Test Connection and Name to test the connection to the client database. You are now ready to copy from the master to the client database.

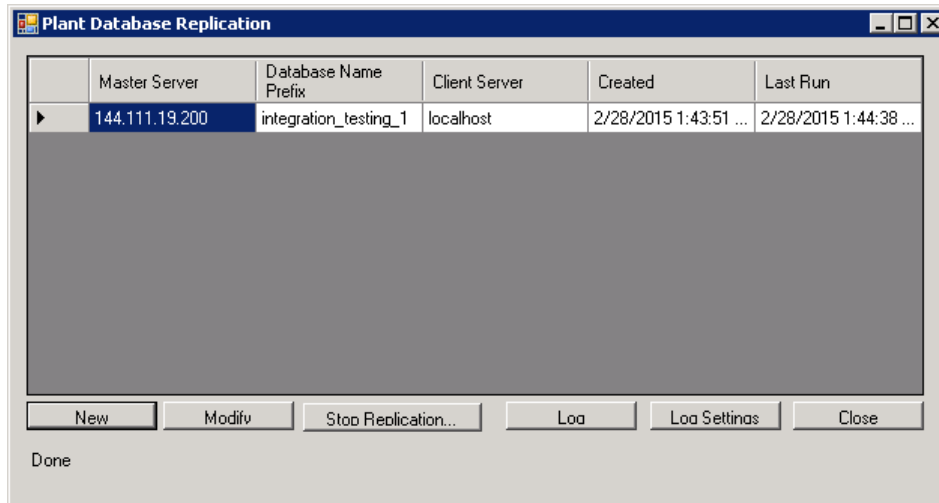
Specify how frequently you want the service to check for updates (applies only when run as a Windows Service as described below). This value is saved in the configuration file used by the synchronization service.

**WARNING - This next step could take quite some time!**

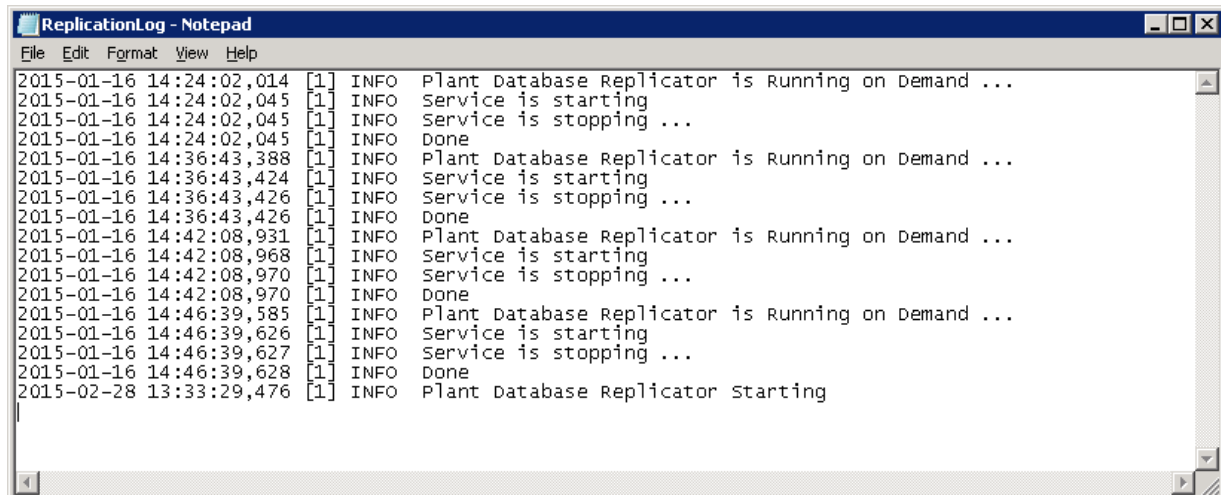
Click Create Client Database to create and copy the database from the master to the client.

Enabling PnPDatabase table data capture

As the database is replicated, status displays in the lower left of the dialog box.



The plant database is now replicated to the client sql server.



Click Log to check the replication log.

The log file and the config file are saved to: *C:\ProgramData\Autodesk\AutoCAD Plant 3D\Replication*

## Synchronization Service

After the plant database is initially replicated in manage mode, you run PlantDatabaseReplication.exe as a service. Running as a SQL Server Job is recommended, but you can also run as a Windows service.

### Run Once

To test synchronization (or if you are using an external scheduler) you can have PlantDatabaseReplication.exe synchronize one time and stop. To run once enter the following at the Windows command prompt:

## PlantDatabaseReplication run

The result of the run will be recorded in the replication log that can be viewed by clicking Log.

### Windows Service


To run replication as a Windows Service it has to be installed first. The installation can be done by using installutils commands:

From the command line enter the following command. Note that you need to enter the full path to PlantDatabaseReplication.exe:

**Installutil c:\Replication\PlantDatabaseReplication.exe.**

This must be run as an Administrator.

After a successful install the PlantDatabaseReplication displays in the Windows Task Manager with the Stopped status:



| Name                     | PID  | Description      | Status  | Group  |
|--------------------------|------|------------------|---------|--------|
| Power                    | 992  | Power            | Running | DcomL  |
| PolicyAgent              | 3400 | IPsec Policy ... | Running | Networ |
| PNRPsvc                  |      | Peer Name ...    | Stopped | LocalS |
| PNRPAutoReg              |      | PNRP Machin...   | Stopped | LocalS |
| PlugPlay                 | 992  | Plug and Play    | Running | DcomL  |
| PlantDatabaseReplication |      | Plant Databa...  | Stopped | N/A    |
| nla                      |      | Performance...   | Stopped | LocalS |

It has to be started in order to perform its work. It can be done from the Windows Task Manager right click menu or by clicking “Start” button in the Plant Replication Manager. The relevant message should be logged into the replication log.

The service can be uninstalled using the following command:

*Installutil /u c:\PlantDbReplication\PlantDatabaseReplication.exe.*

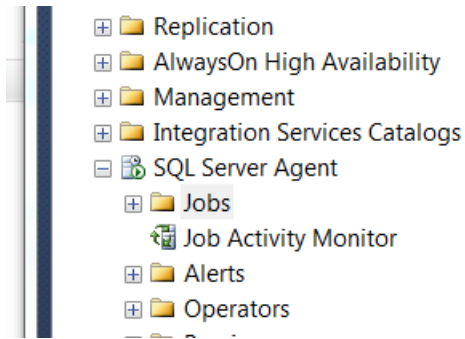
### SQL Server Job (Recommended)

SQL Server Jobs can to be configured from the Microsoft SQL Server Management Studio.

Plant Database Replication modules should be deployed on the machine where SQL Server Instance for which you are planning to add a job is running.



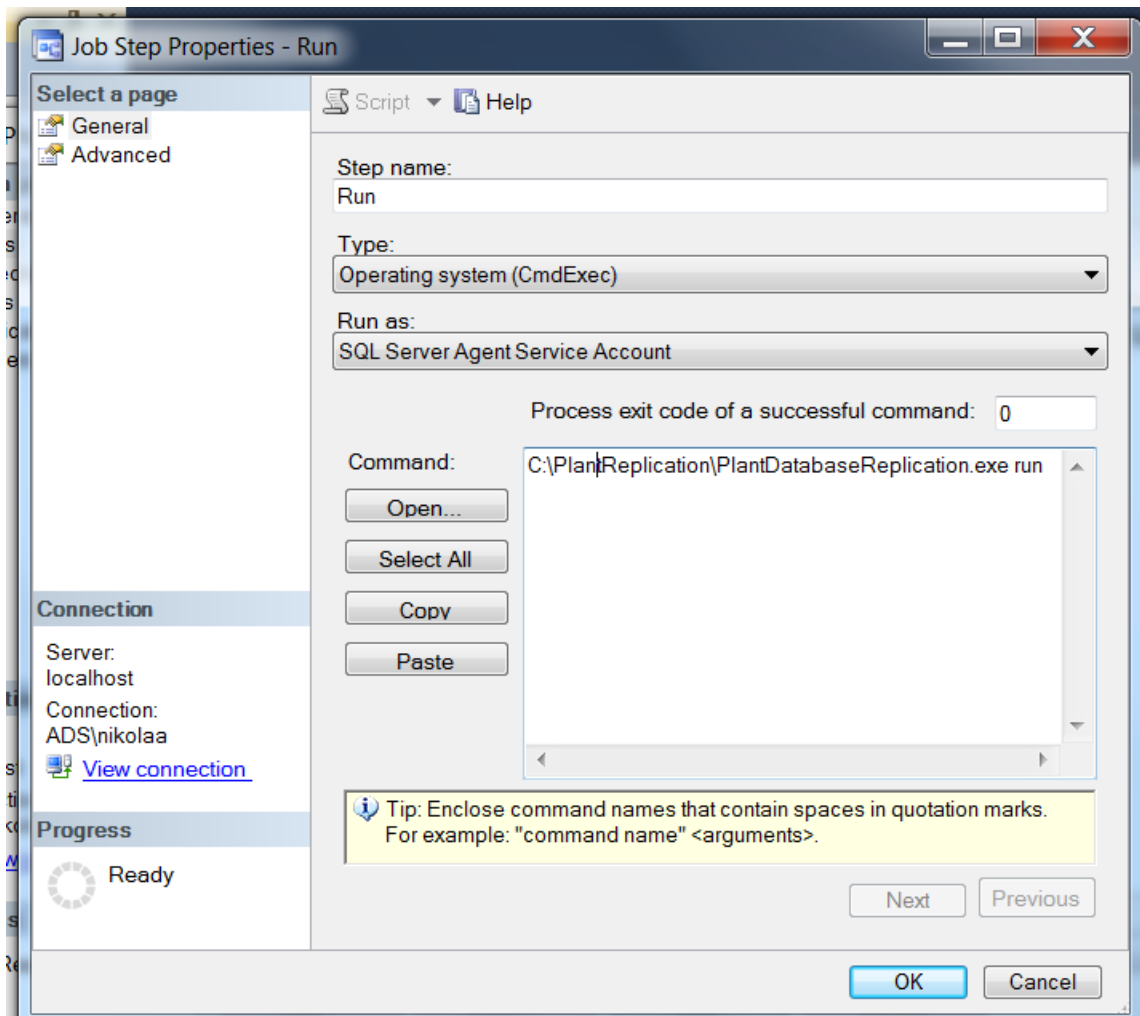
Connect to the SQL Server instance, expand SQL Server Agent node, right click on Jobs and select New Job.



In the General section enter Job name. For instance, Plant Database Replication.

Select “Steps” item in the “Select a page” section at the left. Click “New”. In the New Step Dialog enter step name (“Run” for instance) and chose Operating System in the Type combo. Enter command text similar to the one below and click Ok:

*C:\PlantReplication\PlantDatabaseReplication.exe run*



You may also define job run schedule by selecting Schedule in the “Select a page” section. If schedule is not defined, replication will run only on demand. Schedule can be added / modified later by choosing Properties item from the Plant Replication Job right click menu.

After completion of the steps above, Plant Database Replication job will appear the Jobs sub-tree.

You may run replication on demand by choosing “Start Job at Step...” from the right click menu on a replication job node.

Upon completion the result will be recorded in the replication log.

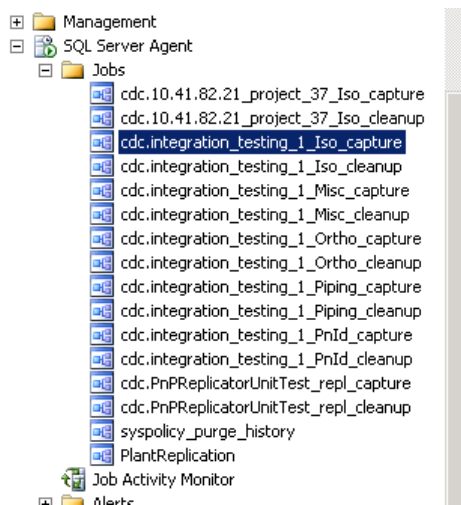
#### Technical Considerations when choosing Windows Service over SQL Server Job

- Windows service process is always running and is keeping some internal state in memory;
- SQL Server job is executed as an operating system command on each invocation and loads its state every time;
- Potentially Windows Service will work faster and will create less network traffic;
- Windows Service can run on any machine on a network that has access both to master and server databases;

- SQL Server Job should be deployed on the machine where client SQL Server instance is running;
- SQL Server Job is tightly linked to the running SQL Server Instance and will start automatically together with the instance;
- Windows Service runs independently from the SQL server instance. It may still run if SQL Server Instance is down, or not run if it is up;

## SQL Server Jobs – Replication Limitation

In SQL Server Management Studio, you can verify that the new plant database has been replicated on the client. You can also see new capture and cleanup jobs.



On SQL Server, there is a limit to the number of jobs that can run.

<http://blogs.msdn.com/b/sqlserverfaq/archive/2012/03/15/inf-limitations-for-sql-agent-when-you-have-many-jobs-running-in-sql-simultaneously.aspx>

This limits the number of projects that can be replicated to 9 (or possibly 10). Because of this, you Admins will want to remove their test databases and shut down replication for inactive projects.


The new CLI **plantvaultdatabase** command makes it easy to stop using the replicated client on the Plant 3D project side. Stop Replication available from **PlantDatabaseReplication.exe** (Manage Mode) removes the jobs from the server.

If you exceed the limitation, you will see the following error in the log file:

```
FATAL GetCDCSessionTime not found after 200 tries 2015-01-14 14:22:09,014 [7]
FATAL select max (end_time) from sys.dm_cdc_log_scan_sessions where end_time >= '2015-01-14
14:21:58.75'
```

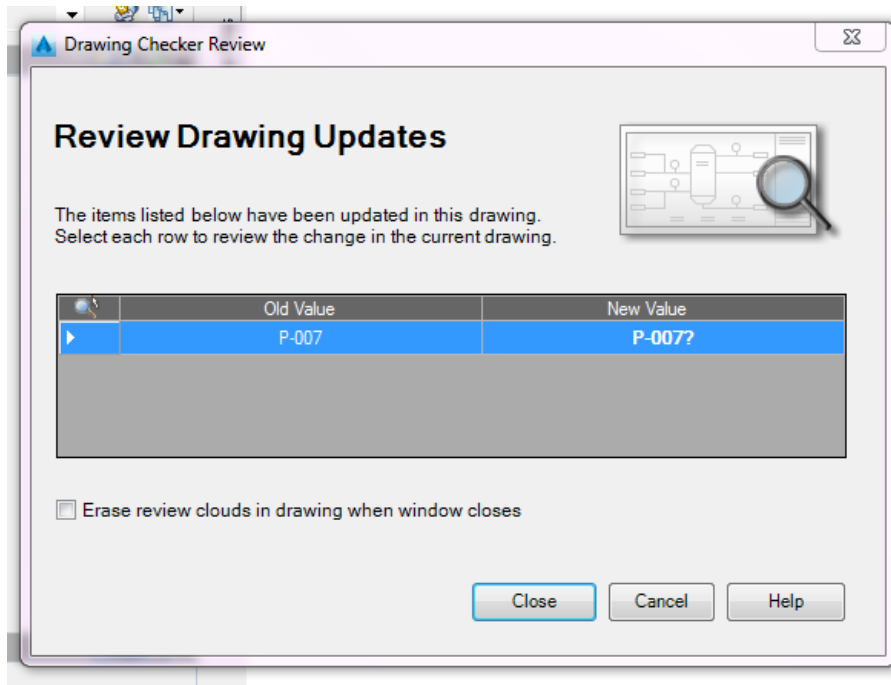
## Replication Topology and Conflicts

Plant Database Replication performs synchronization between two databases – master (usually the default plant database) and client. Client database can be connected only to one master. Master may have multiple client databases referencing it. Client database keeps information about master it is connected to.



In the picture above A is a master (and would usually be the default) database; B and C are its clients. B is a master for D and E. There are no limitations on how many databases can be replicated or how they are nested, however conflicts and delays can be avoided with fewer replicated servers.

**All possible conflicts that may happen during synchronization are resolved in the Master favor.** For instance, if the same row was modified by master and client, master will overwrite client changes (resolution is done on a row basis, not individual properties).



The Drawing Checker Review dialog displays when you re-open a P&ID drawing following a tag conflict.

If two users use the same tag and check in the file at the same time to different vaults, the client database tag becomes unassigned, and displays a question mark in its tag.

Conflicts are identified in the LOG file.

## Disabling Replication in SQL Server

You can remove the capture jobs and drop synch tables in when you no longer want to replicate the plant database in SQL Server Management Studio or using sqlcmd.

```
DisableDataCapture.sql:
/*****
/* Disable change tracking */
*****/
use $(dbname)
declare
@enabled int,
@db varchar (155)
set @db='$(dbname)'
set @enabled = (select is_cdc_enabled from sys.databases where name=@db)

if @enabled = 0
begin
    print 'Data capture for database "$(dbname)" is already disabled'
end
else
```

```
begin
    /* Disable database */
    exec sys.sp_cdc_disable_db
    print 'Data capture for database "$(dbname)" has been disabled'
end
```

*DropSynchTable.sql:*

```
/* **** */
/* Disable change tracking */
/* **** */
use $(dbname)
go

if EXISTS (SELECT * FROM sys.objects WHERE object_id =
OBJECT_ID(N'[dbo].[SyncDatabase]') AND type in (N'U'))
begin
    Drop table [dbo].[SyncDatabase]
end
```

1. Create and run *DisableDataCapture.sql* and *DropSynchTable.sql* scripts (text above) for the database of interest. For example:

```
sqlcmd -S <servername> -E -i DisableDataCapture.sql -v dbname="<database to disable>"
sqlcmd -S ccs2ua0510y51 -E -i DisableDataCapture.sql -v dbname="PnPReplicatorUnitTest_repl"
sqlcmd -S ccs2ua0510y51 -E -i DropSynchTable.sql -v dbname="PnPReplicatorUnitTest_repl"
```

2. Remove corresponding database entry from the rconfig.xml file

## Maintenance and Manual Audit

If you get a *Metadata not in site* error message the plant database may need to be manually reconciled.

**PLANTVAULTRECONCILEPROJECT** updates sync related data in the master database.

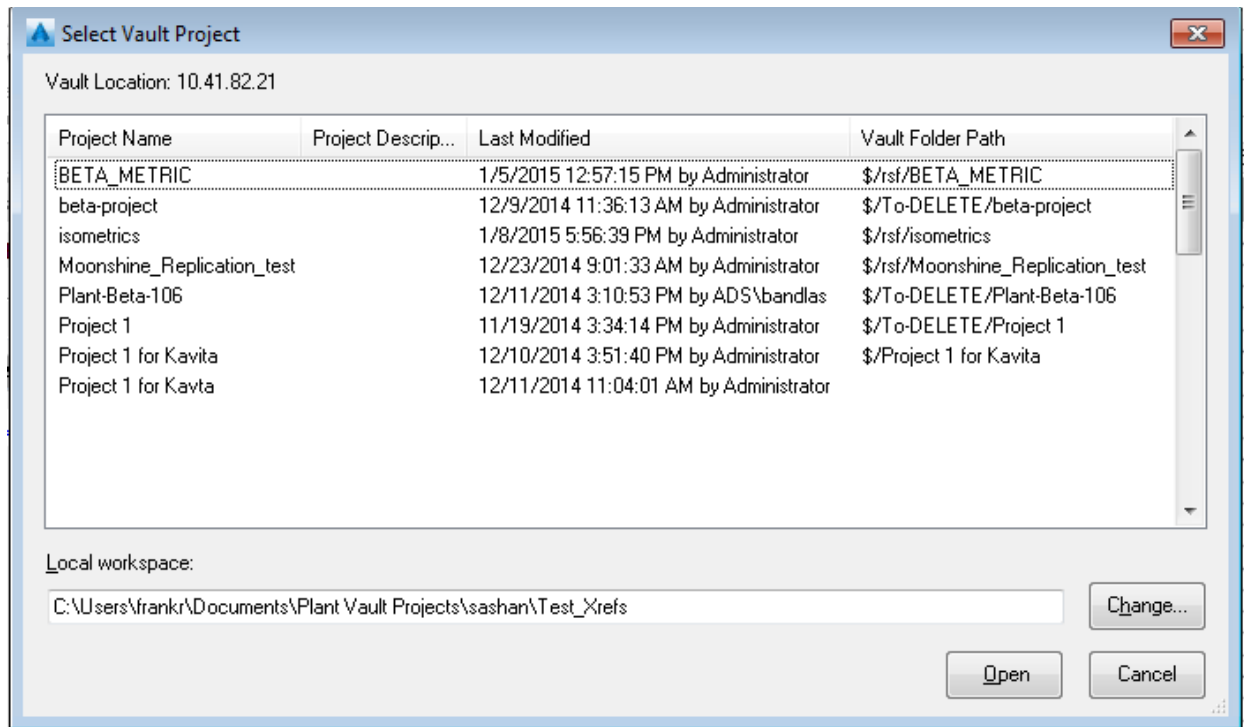
If the log files contains *GETCDCSessionTime* fatal error do the following:

- Identify the capture job that is not running (it should be in the database that is associated with Vault server that exhibits the error (the replication log does not have info about the database where error occurred). Job names look like cdc.<project database prefix>\_<part suffix>\_capture (for example: cdc.ccs1pd3p12\_project\_rtm\_test\_Ortho\_capture).
- Start the job by executing the sys.sp\_cdc\_start\_job stored procedure
- In AutoCAD Plant 3D, open the project from Vault server that displays the *Not on site* error message and run **PLANTVAULTRECONCILEPROJECT** at the Command prompt.

## Project Setup - PLANTVAULTDATABASE

After replication is running, use the **plantvaultdatabase** command to specify which plant database to connect to when you open a project from a specific vault.

This command is run from AutoCAD Plant 3D or AutoCAD P&ID at the Command: prompt.



This command is documented in product help: [PLANTVAULTDATABASE](#).