ThorApp's

No nonsense tools, apps and add-ons for SharePoint

Project Integration Engine 2016 Configuration Guide

Author: Michael Bear
Date: 28 Sep 2018
Version: 1.1

Overview

This document describes the steps to configure Project Integration Engine (PIE) 2016 for Project Server 2016. After reading this guide you should know how to configure the Project Integration Engine 2016 Site collection feature, including:

- Creating Data Sources
- Creating Datasets
- Creating Processing Packages
- Creating Processing Schedules

Prerequisites

To follow this guide, would have completed the Installation and Deployment Guides, and have the following items:

- Administrator rights in Project Server
- Farm Administrator rights in SharePoint Central Administration
- A basic understanding of Project Server Internal Field Names

Data Sources

First, you need to create a Data Source.

Navigate to Project Integration Engine home page at http://[site collection]/pwa/pie
Click "Data Sources" in the Quick Launch.

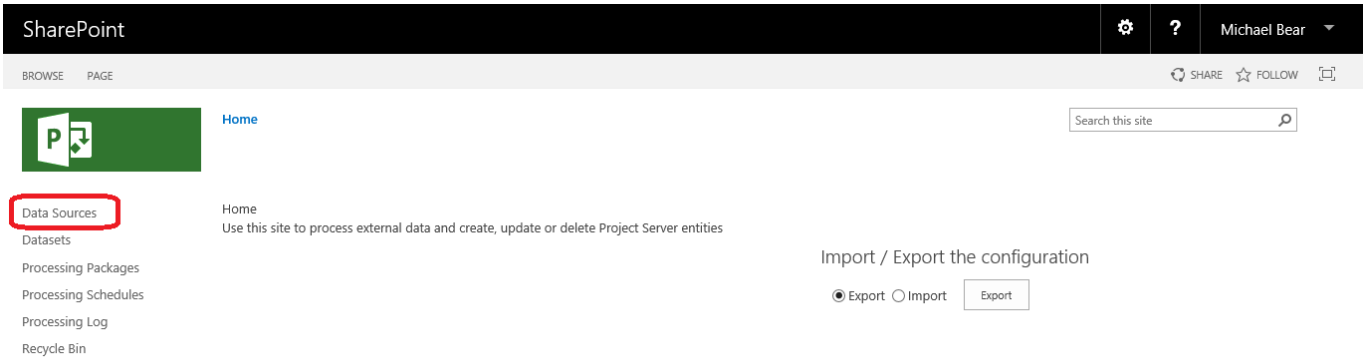

Click "New Item".

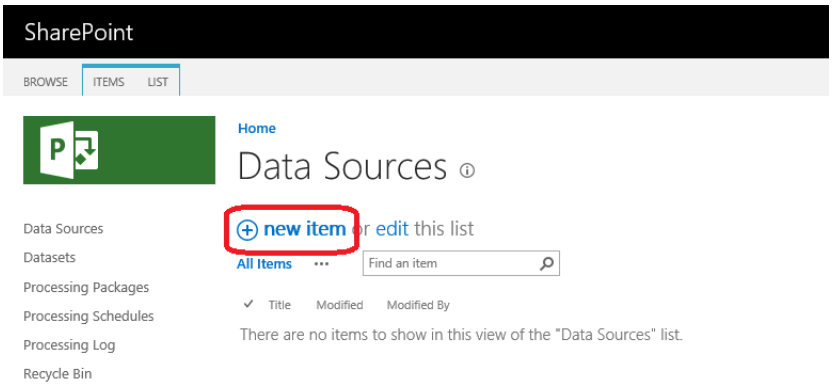

Enter a name of the new Data Source.

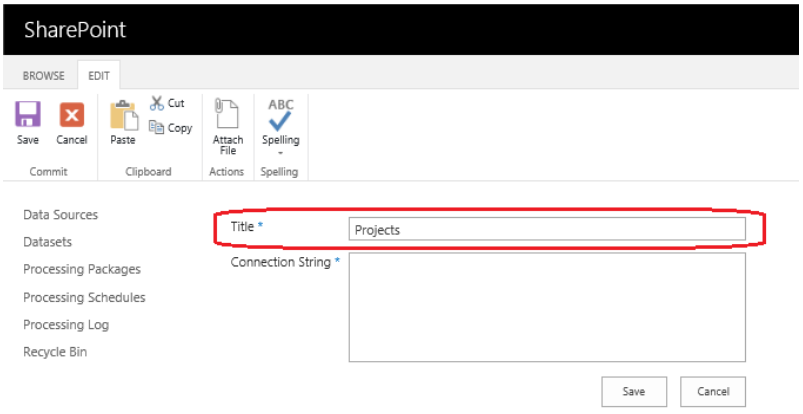

To get the Connection String, on the desktop, create a new .txt file.
Rename it to 'something.udl'

Open this file.
Enter the database name that contains the data you're after.
Tick the Allow Saving Password check box.
Enter user credentials.
Click "OK".

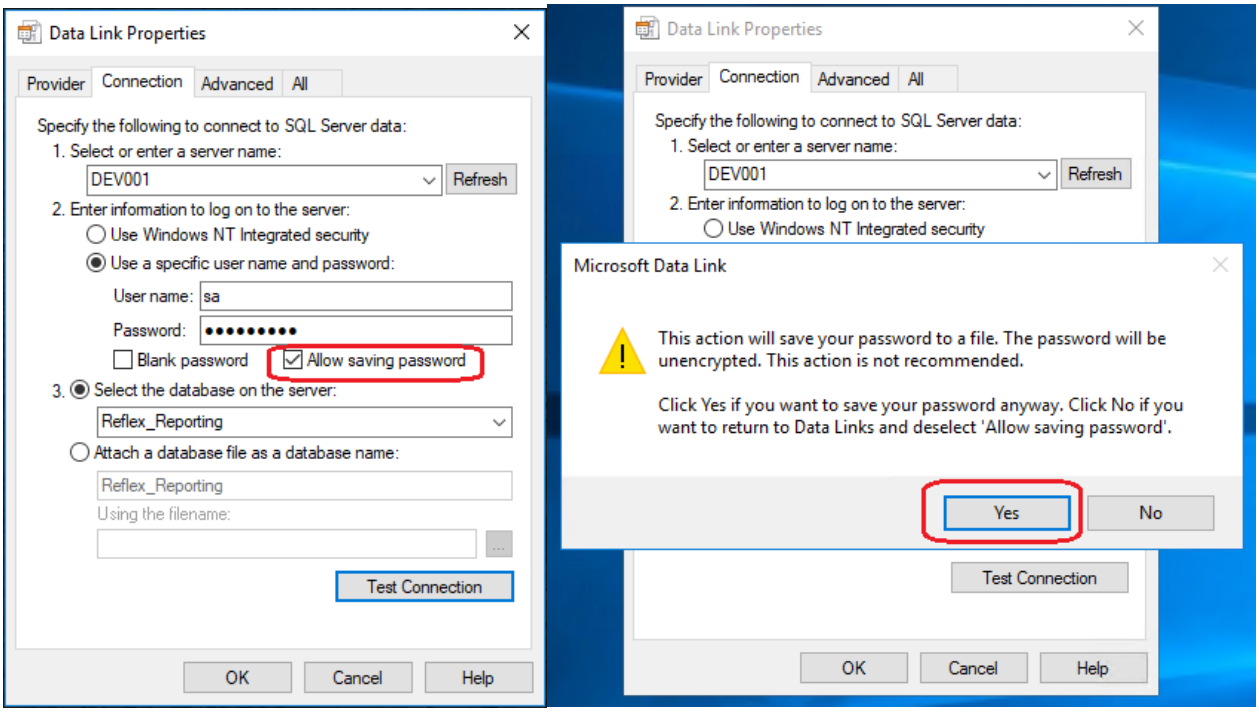

Rename to 'something.txt'.
Open the file in Notepad.

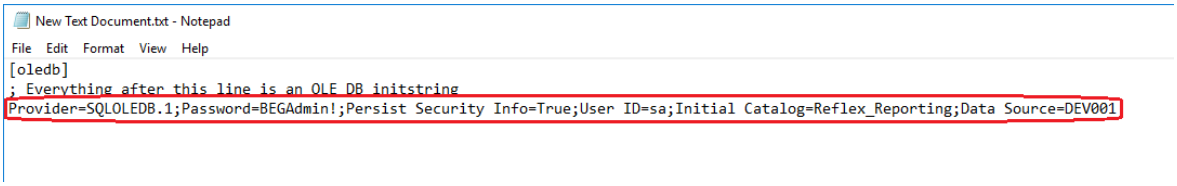

Highlight and copy the text indicated in red.
Return to the browser and paste the copied text into the Connection String text box.
Click "Save".

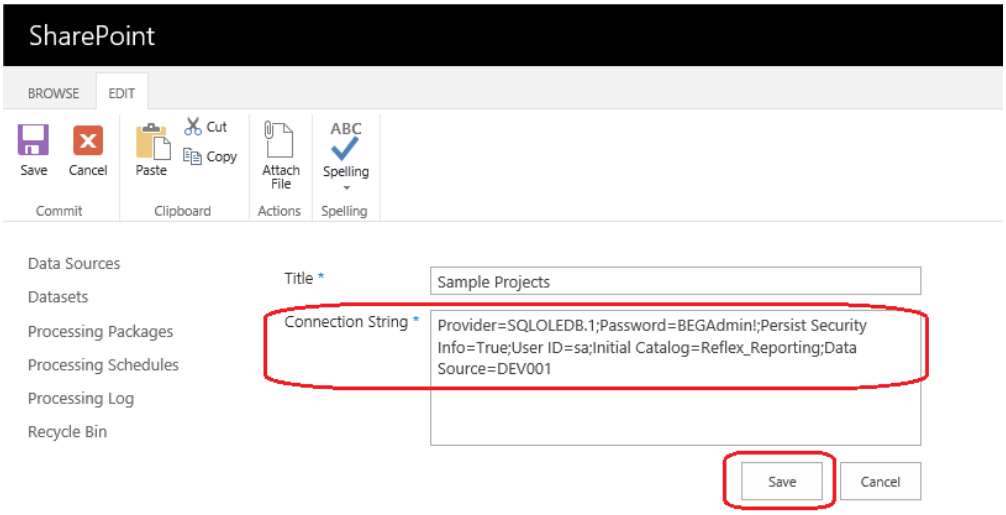

You will now see the new Data Source listed.

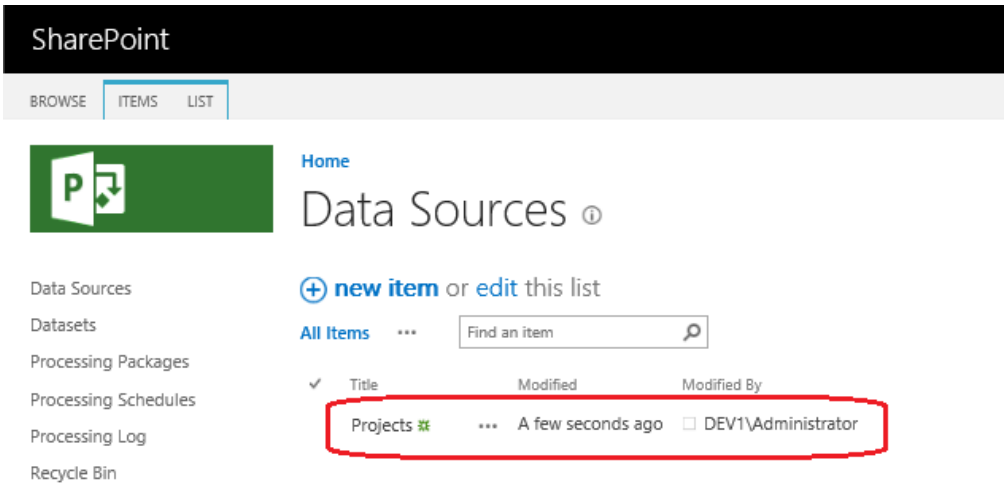

Datasets

Next, you need to create a Dataset.

Click on "Datasets" from the Quick Launch.
Click "New Item".

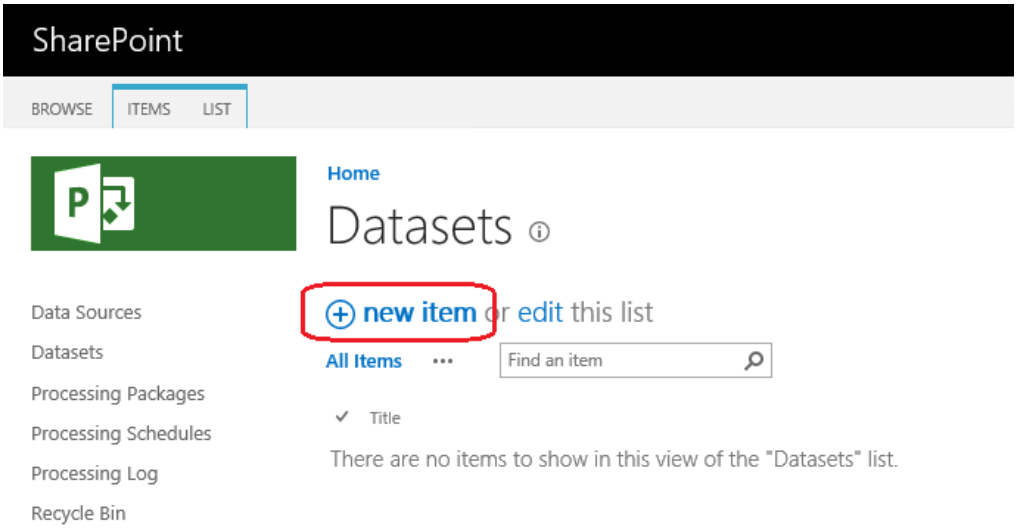

Enter a name for your Dataset in the Title textbox.
Select the Data Source from the drop-down-list.

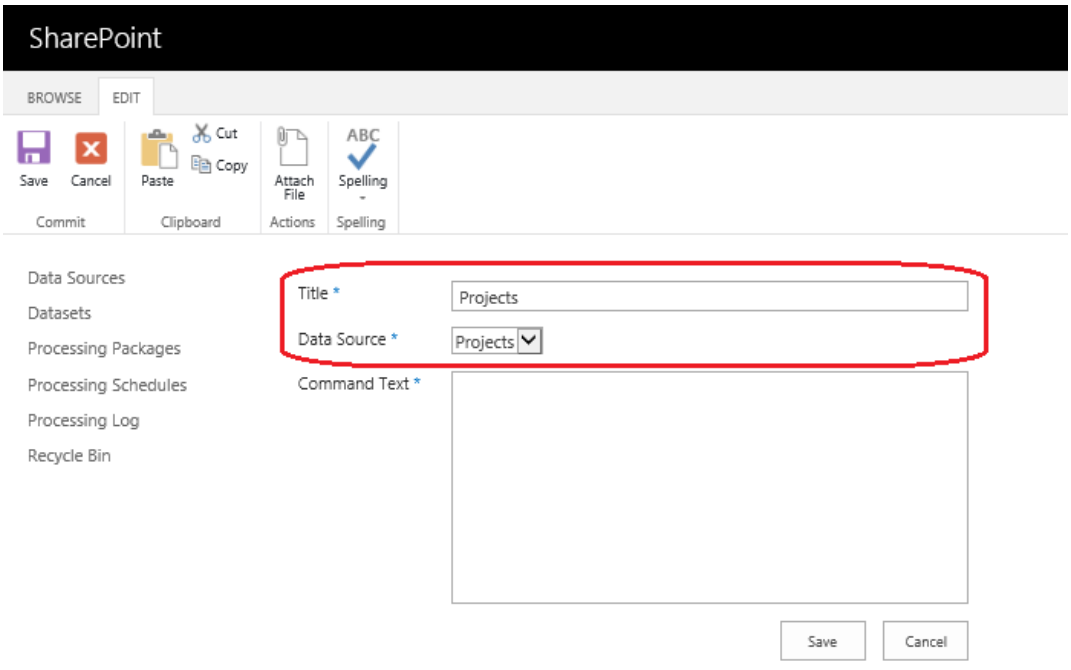

The Command Text is a SQL query that will fetch data from a table or view you want to import into Project Server. You can generate the Command text from an ODC connection using Excel.

Open Excel in a new blank/new workbook.
Select "Data" from the ribbon bar.
Select "Get Data".
Select "From Other Sources".
Select "From OLEDB".

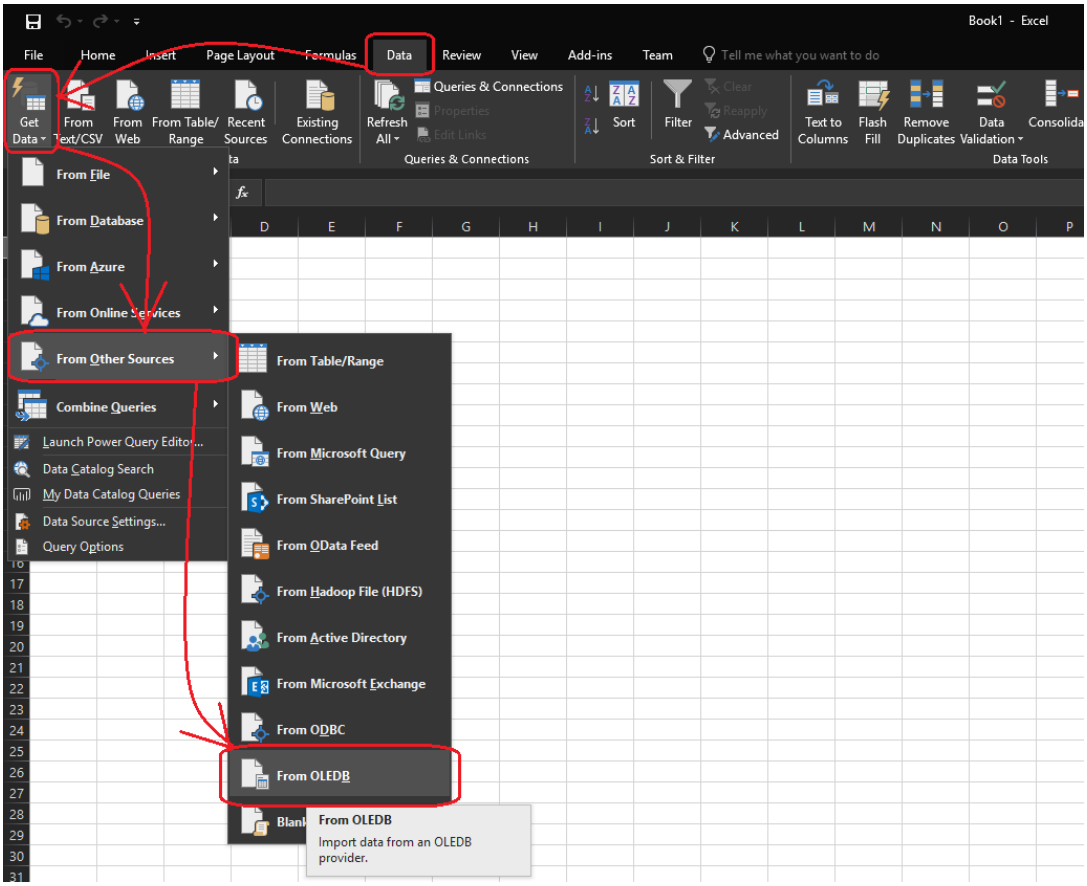

You will be presented with the following dialogue box.

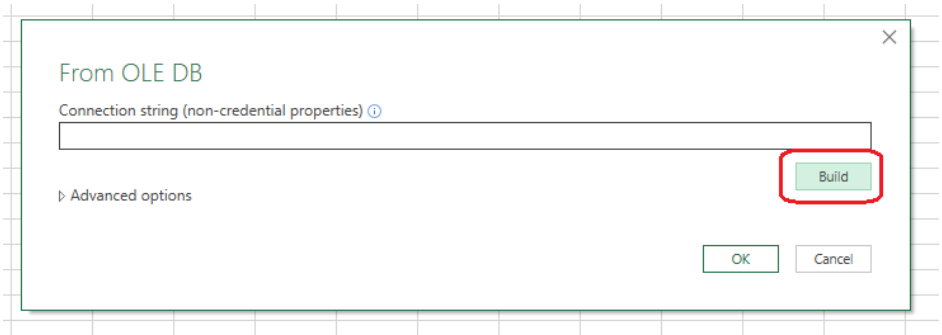

Click "Build".
In the Data Link Properties window, select SQL Server Native Client.
Click on the "Connection" tab.
Enter the Server Name, Database name and credentials.

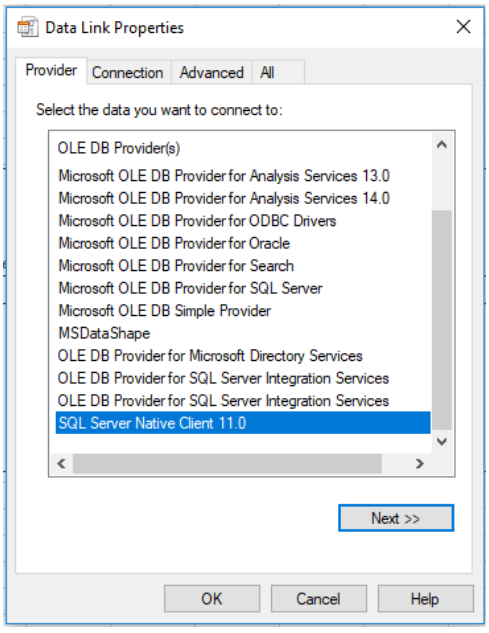

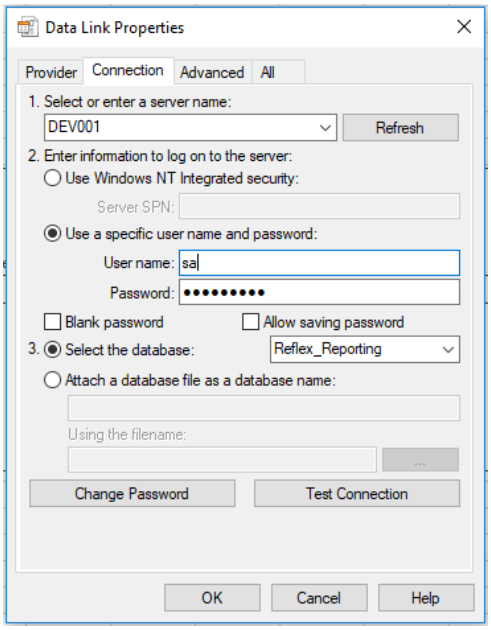

Click "Test Connection" to confirm a valid connection.
Click "OK".
Click "OK" to generate the connection string.

Click "OK".
You will be prompted for a user name and password, and optional "Credential connection string properties". Enter the appropriate credentials and click "Connect".

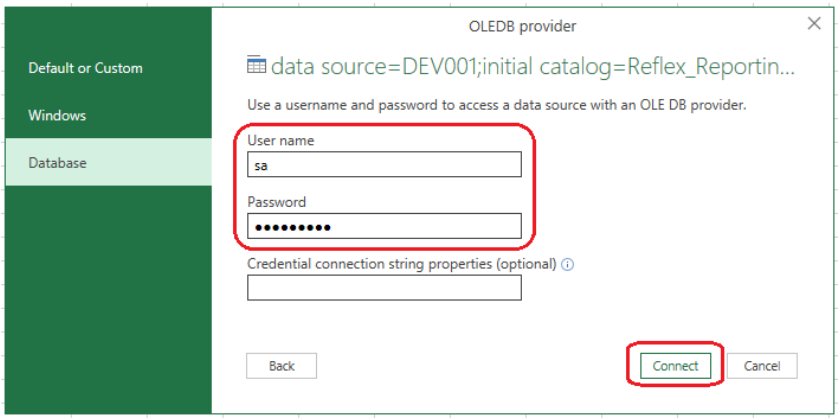

In the Navigator window, select the table required:

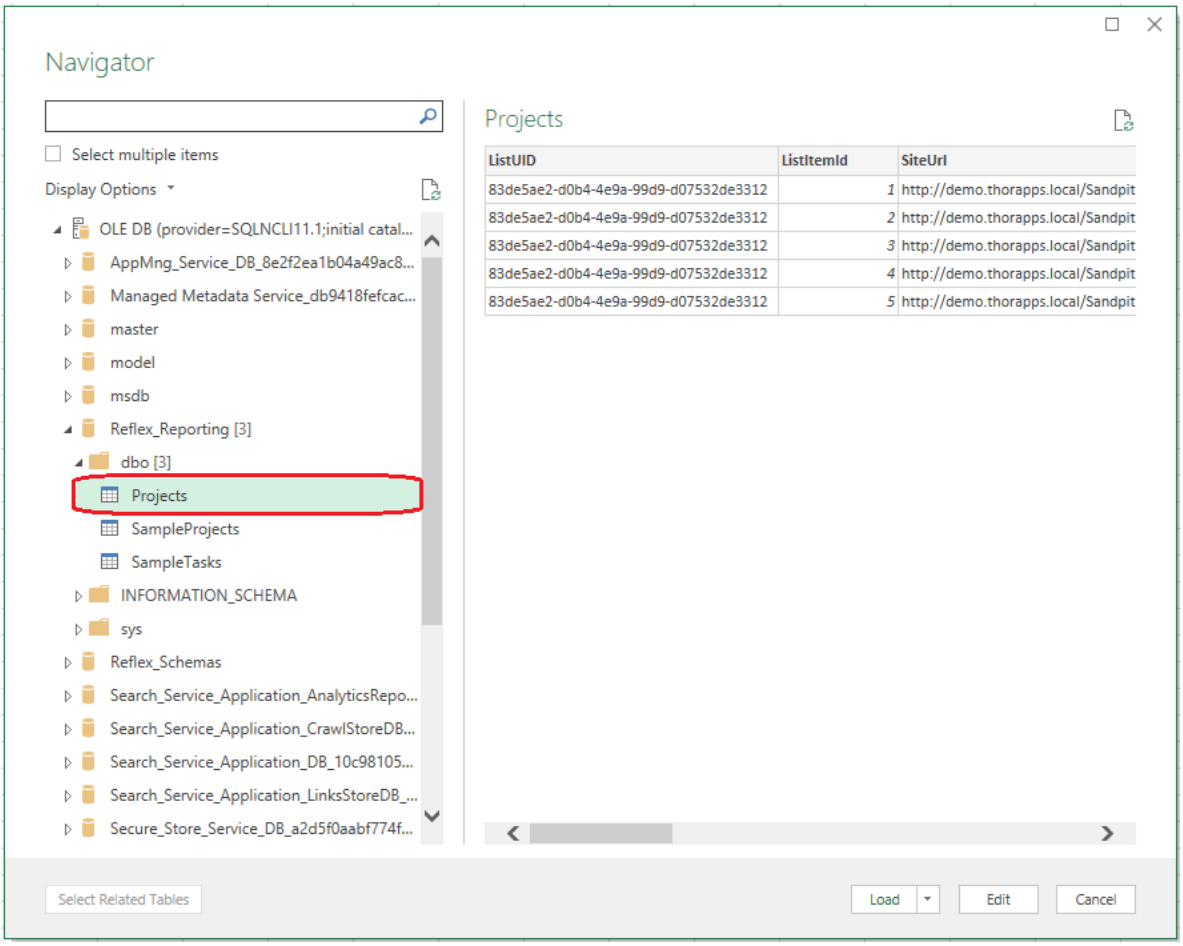

The table data will be displayed in the workbook.

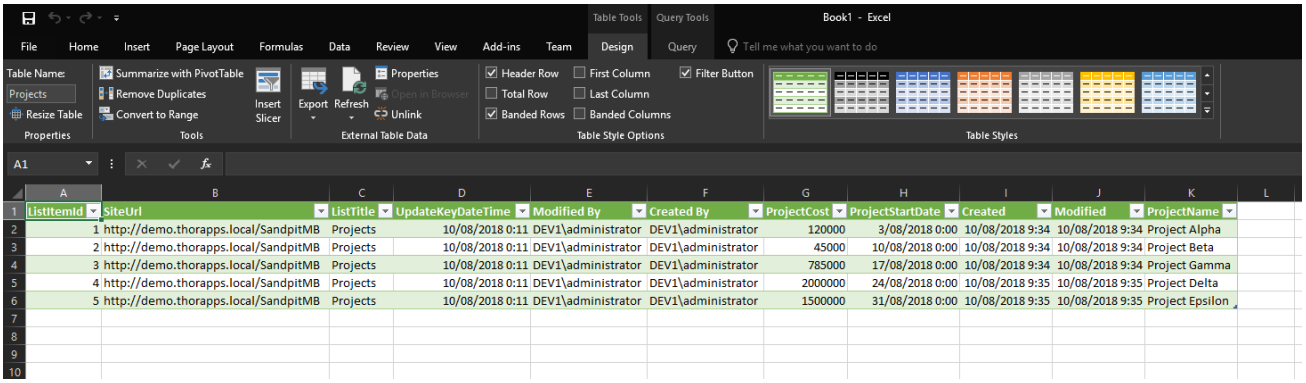

Now, we need to export the ODC file.

Under the Data tab, click "Queries & Connections".
In the pane that opens, click "Queries"
Right click on the query displayed.
Select "Export Connection File...".

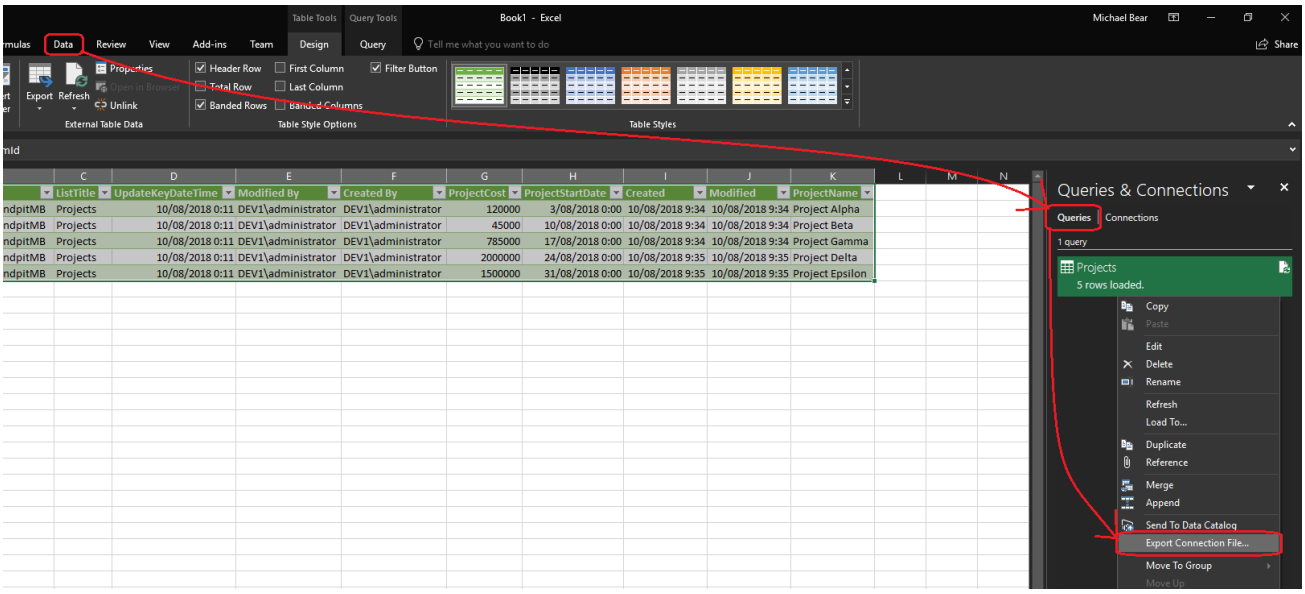

Give it a name, and save it as an ".odc" to the desktop (or other convenient location).
Once saved, locate the file, rename it as a ".txt" file and open it with notepad.
We need the text located within the "<odc:CommandText> </odc:CommandText>", as shown highlighted below. Copy this text to the clipboard.

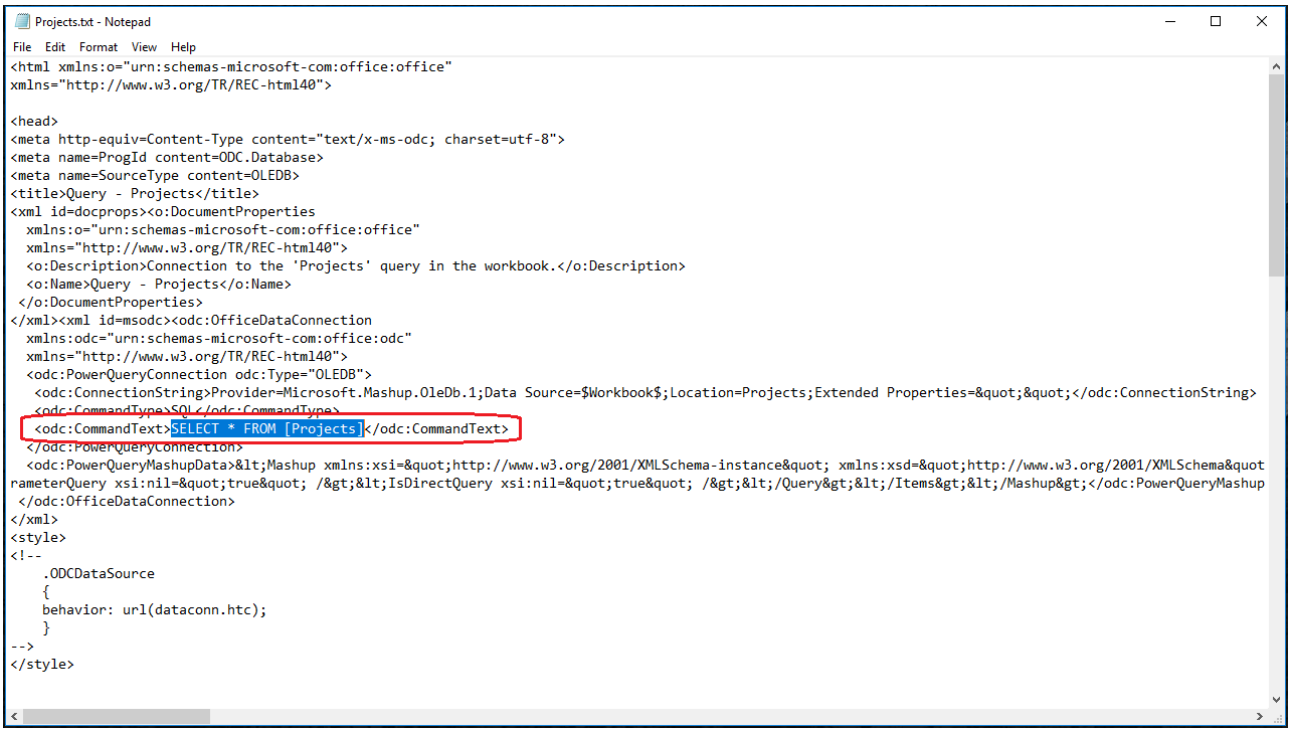

Return to the browser, and paste this text into the "Command Text" text box.
Click "Save".

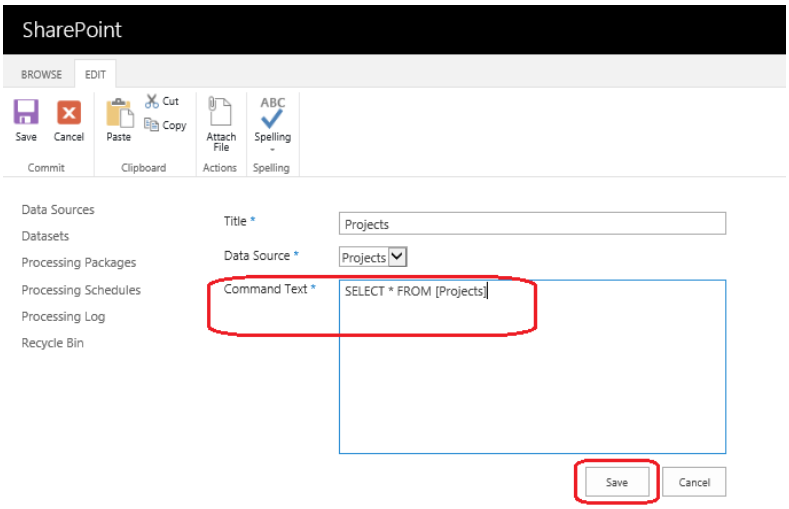

The new Dataset is now displayed in the list:

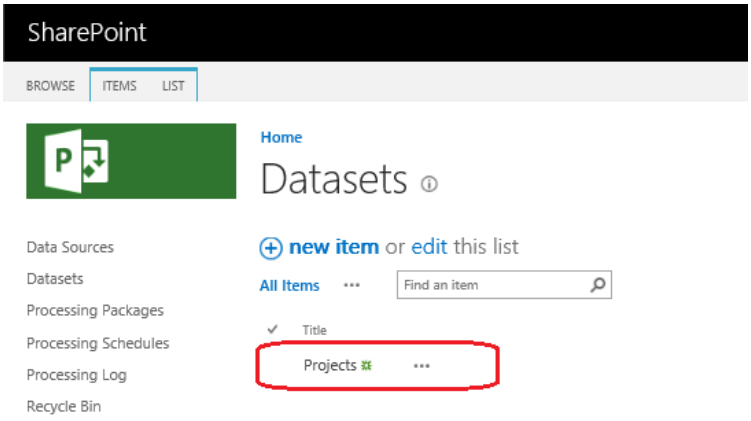

Processing Packages

A Processing Package defines how the Project Integration Engine will process each row from a Data Source and how it relates each row to existing content.

Step 1

Select "Processing Packages" from the quick launch.
Click "New Item".

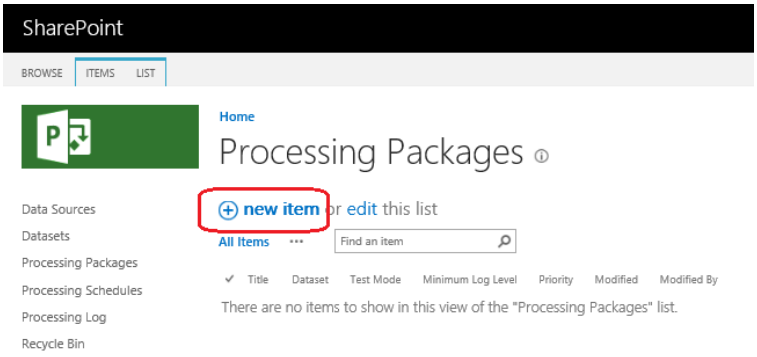

Enter a value for Title and select your Data Source.

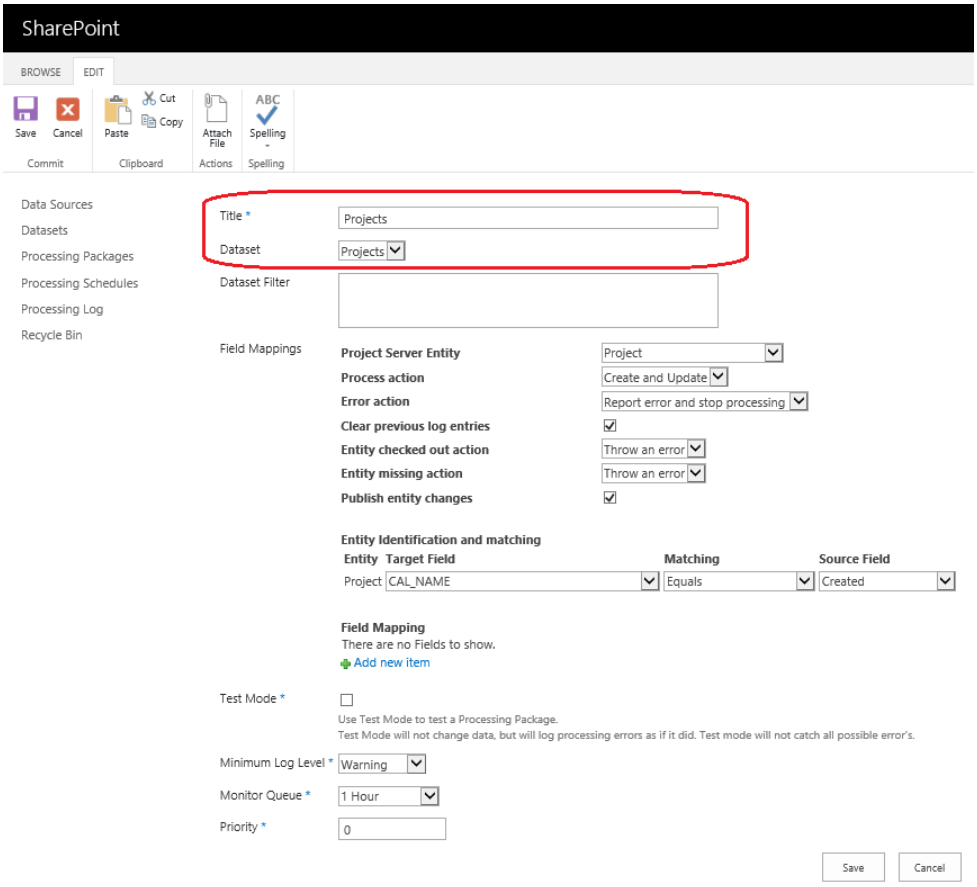

Wait for the page post-back to complete as it needs to query the data source to get the list of source fields to display. This may take a while depending on your Data Source query design.

Step 2

Check that…

- "Project Server Entity" is set to "Project"
- "Process action" is set to "Create and Update"
- "Error action" is set to "Report error and keep processing"
- "Clear previous log entries" is checked
- "Entity checked out action" is set to "Force check-in"
- "Project missing action" is set to "Throw an error"
- "Publish entity changes" is checked

Under "Entity Identification and matching" set

- Target Field to "PROJ_NAME"
- Matching to "Equals"
- Source Field to "ProjectName"

Under Field Mapping create 3 new rows with the following values. Be aware, you need to save each row as it is added, before adding a new row.

Target Field	Source Field	Conversion Rule
PROJ_NAME	ProjectName	Apply Value
PROJ_INFO_START_DATE	ProjectStartDate	Apply Value
Sample Proposal Cost	ProjectCost	Apply Value

*Ignore the "Value Delimiter" field. *This mapping is NOT a direct field for field mapping

Step 3

Leave the remaining fields at their default values and click “Save”.

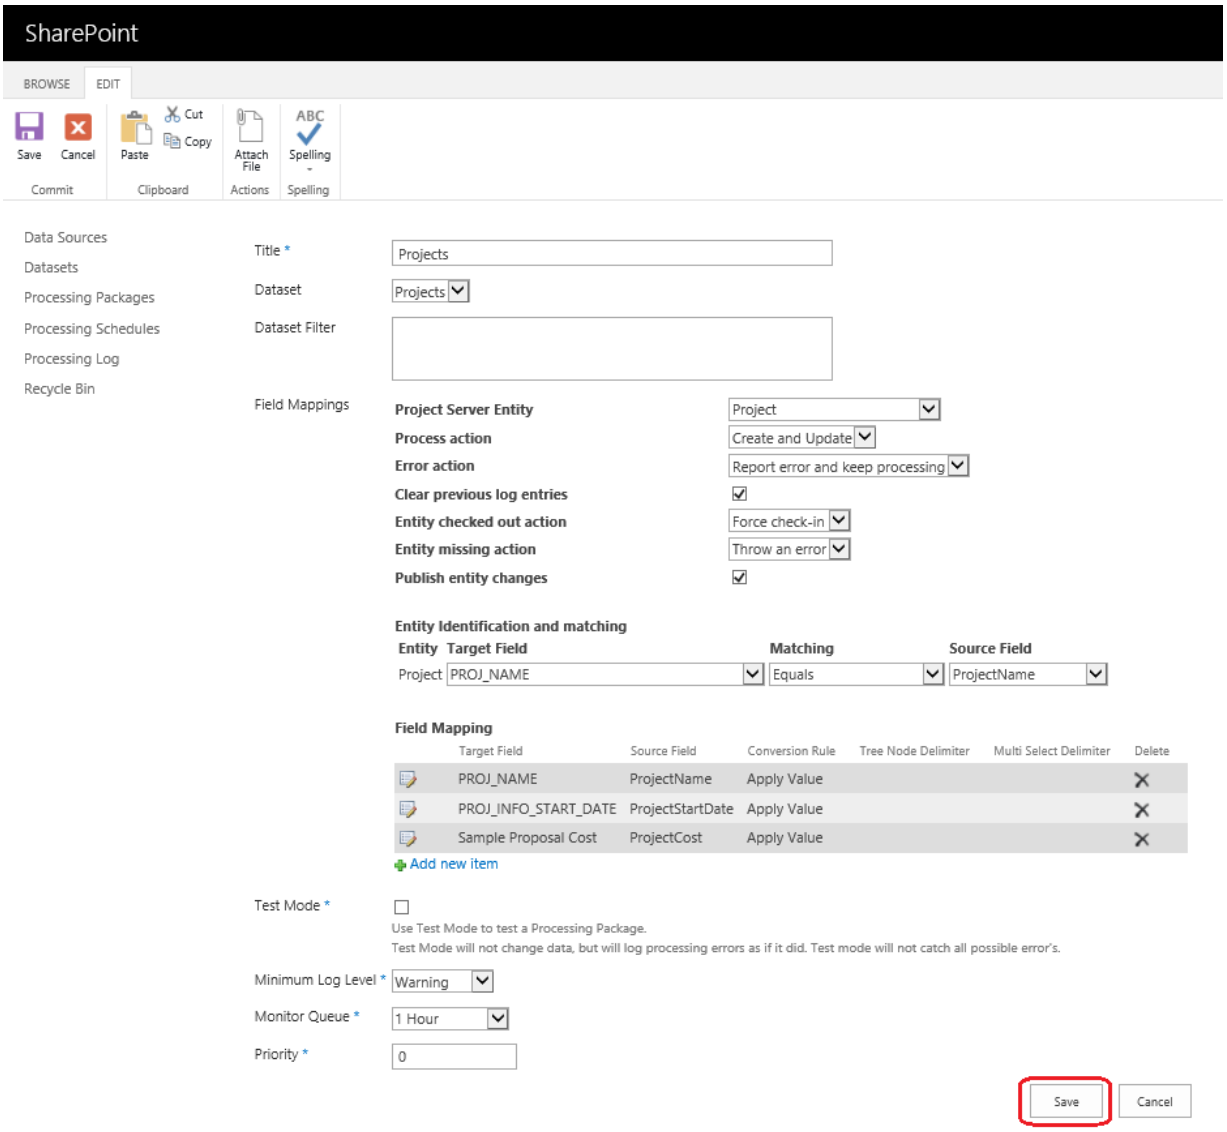

This will create a Processing Package. The next task is to create a Processing Schedule.

Notes on Processing Packages

The Processing Package item page will change depending on the "Project Server Entity" value selected under "Field Mappings". The section below gives a brief overview of the page areas and purpose.

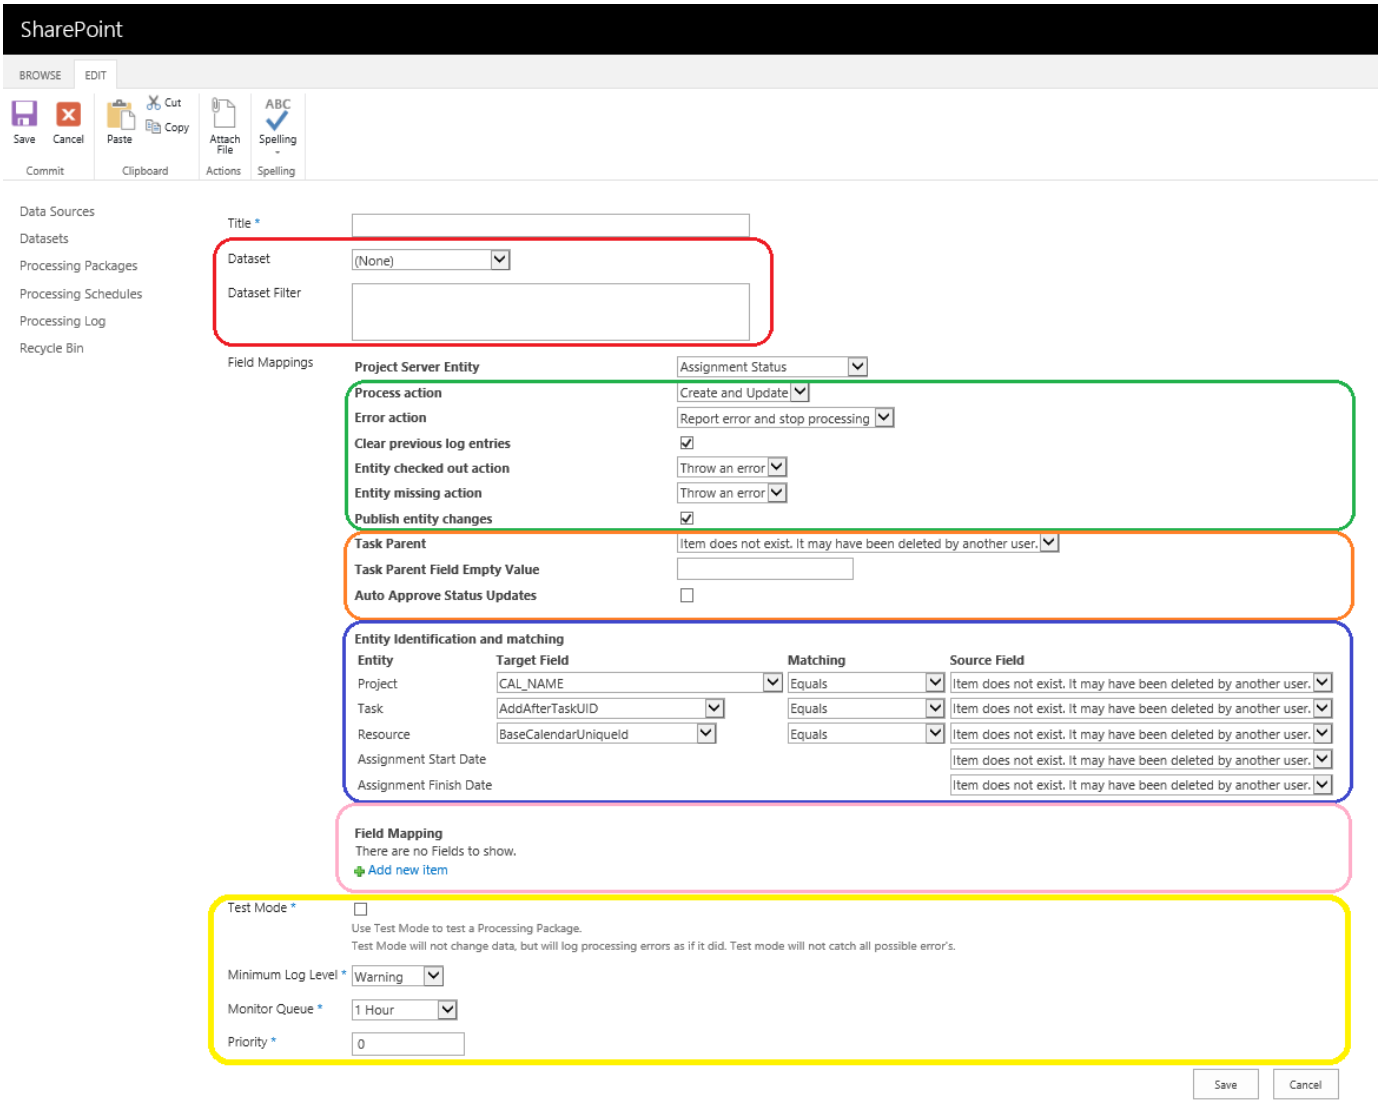

The Red area doesn't change and is used to select and filter the source data.
The Green area doesn't change and controls the actions to take when processing.
The Orange area changes depending on the "Project Server Entity" selected and defines Entity specific processing variables.
The Blue area changes depending on the "Project Server Entity" selected and defines how rows from the Data Source are related to existing content in Project Server.
The Pink area is changed by the user and controls which and how target fields are populated from source fields.
The Yellow area doesn't change and controls how to process and what to log.

Processing Packages - Detailed Field Description

Field Name	Value/s	Description
Title	User defined	The name of the Processing Package
Dataset	Selected from the list of available Datasets	The Source Data used by the Processing Package. The Source Data must be singular rows of a single Project Server Entity type (i.e. all Projects and no Tasks, or all Tasks and no projects etc.)
Dataset Filter	User defined	The Dataset filter is a simple filter applied to the set of rows after they have been retrieved from the Dataset. This is an inefficient method of filtering. Preferably the data should be filtered by the source system.
Project Server Entity	Project Task Resource Resource Rate Resource Calendar Exception Assignment Assignment Status Lookup Table	This field determines what Entity type is being processed and therefore how the row data is processed. Changing this field causes a page post back and changes the fields available on the page.
Process action	Create and Update Create only Update only Delete	Determines what action you want the processor to apply to Project Server
Error action	Report error and stop processing Report error and keep processing	Determines what the processor should do if an error is encountered while processing a Data Source row
Clear previous log entries	Checked Unchecked	Deletes old entries from the Processing Log list for this package when the package processing is started.
Entity checked out action	Throw an error Force check-in Skip item	Determines what action the processor should take if the entity to process is checked out
Project missing action	Throw an error Skip item	Determines what action the processor should take if the Project for an entity is missing. This field only applies to Projects, Tasks and Assignments.
Publish entity changes	Checked Unchecked	Determines if the processor should publish changes after processing a Project, Task or Assignment row
Task Parent	Selected from the field names available on the Data Source	Tasks can exist in a hierarchical structure, but the Dataset only allows for flat table rows. Therefore, to define the relationship between parent and child tasks there must be a field in the Dataset row that determines which other row is the parent task.

Task Parent Field Empty Value	User defined	If the Dataset contains a hierarchical structure of Tasks and a field is defined as the Task Parent field, there must be a value for Tasks with no parent (i.e. level 0 tasks)
Auto Approve Status Updates	Checked Unchecked	If processing Assignment Status rows, checking this field will also Approve the changes made to the Status Updates.
Lookup Table Value Tree Delimiter	User defined	Lookup tables can be a hierarchical structure, but the Dataset only allows for flat structures. Usually a lookup tree value is defined by a delimited list of values defining the values position in the tree. This field allows you to define the delimiter character (or string) used in the Dataset
Test Mode	Checked Unchecked	Determines if the changes should actually be made to Project Server or just tested. Checking this field will apply no change to Project Server, but will not catch all possible errors
Minimum Log Level	Information Warning Error	Determines the minimum event level to log in the Processing Log list. Logging all Information items can add significant load to the server. It is recommended this value be kept at Warning.
Monitor Queue	Don't Monitor [Some Times] Indefinitely	The Project Integration Engine uses the Project Server queue to process changes into Project Server. By default it will keep track of each change and won't start the next change within the same session until the first has completed (or errored). Setting this value to "Don't Monitor" will just log the jobs into the Project Server queue and move on to the next task, however if an error was to occur, the Project Integration Engine will not know about it and continue to process without logging the error. Recommended setting for this field is "1 Hour". This leaves plenty of time for the Project Server queue to catch up, without causing the processor to fail.
Priority	User defined integer	This field is used to determine the order of execution of Processing Packages when bundled together in a Processing Schedule. Processing Packages are executed in ascending Priority order (i.e. 0 is first, followed by 1 and then by 2)

Entity Identification and matching

The Project Integration Engine assumes the Dataset contains a unique set of rows for the Project Server Entity type selected. For example, if the Project Server Entity "Project" is selected, the Dataset must contain only 1 row for each project. If the Project Server Entity "Assignment" is selected the Dataset must contain only 1 row for each Project, Task and Resource combination. Therefore, the Project Integration Engine needs to know which fields determine the uniqueness of a row in the Dataset. When processing a Dataset that contains rows of Project Server Entity "Project" this can be determined with a single field

(usually a Project Number or Name). When processing a Dataset that contains rows of Project Server Entity "Task" this needs to be determined with 2 fields (1 to identify the Project the task relates to and the second to identify the Task itself)

In addition to identifying the fields that determine the row uniqueness, the Project Integration Engine also needs to know how each of these fields map to a field for the entity within Project Server. This information is used to determine if the row should result in adding new data, updating existing data or deleting existing data in Project Server.

The Entity Identification and matching field "Matching" should always be "Equals" and only changed by advanced users that fully understand the meaning of the alternatives.

Field Mapping

The Field Mapping section allows you to determine which field in the Data Source is applied to which field (for the Project Server Entity selected) in Project Server. The Target Field is the Project Server field being populated and the Source Field is the field from the Data Source being read.

Obviously, the Field Mapping section does not apply when "Process action" is "Delete" and can contain no field mapping rows.

The Conversion Rule field allows you to either simply ignore this field mapping, directly apply the value from the Source Field, or apply the value from a lookup table where the lookup table value description contains the information that can be seen in the Source Field.

The Value Delimiter field is used when the Source field is a delimited list of values defining a node in a tree structure. This is used when the Target field is a Project Server custom field that references a Lookup table that has a treed structure and contains duplicate tree node values but in different tree branches (i.e. Department.Branch 1.John Smith is different to Department.Branch 2.John Smith).

Processing Schedules

Processing Schedules are used to define a set of Processing Packages to process at a predetermined time and frequency. Processing Schedules are picked up by the SharePoint Timer job Project Integration Engine scheduled job processor, and executed within 5 minutes of the Next Start Date.

Processing Packages selected in the Processing Schedule are executed in an order determined by the "Priority" field on the Processing Package.

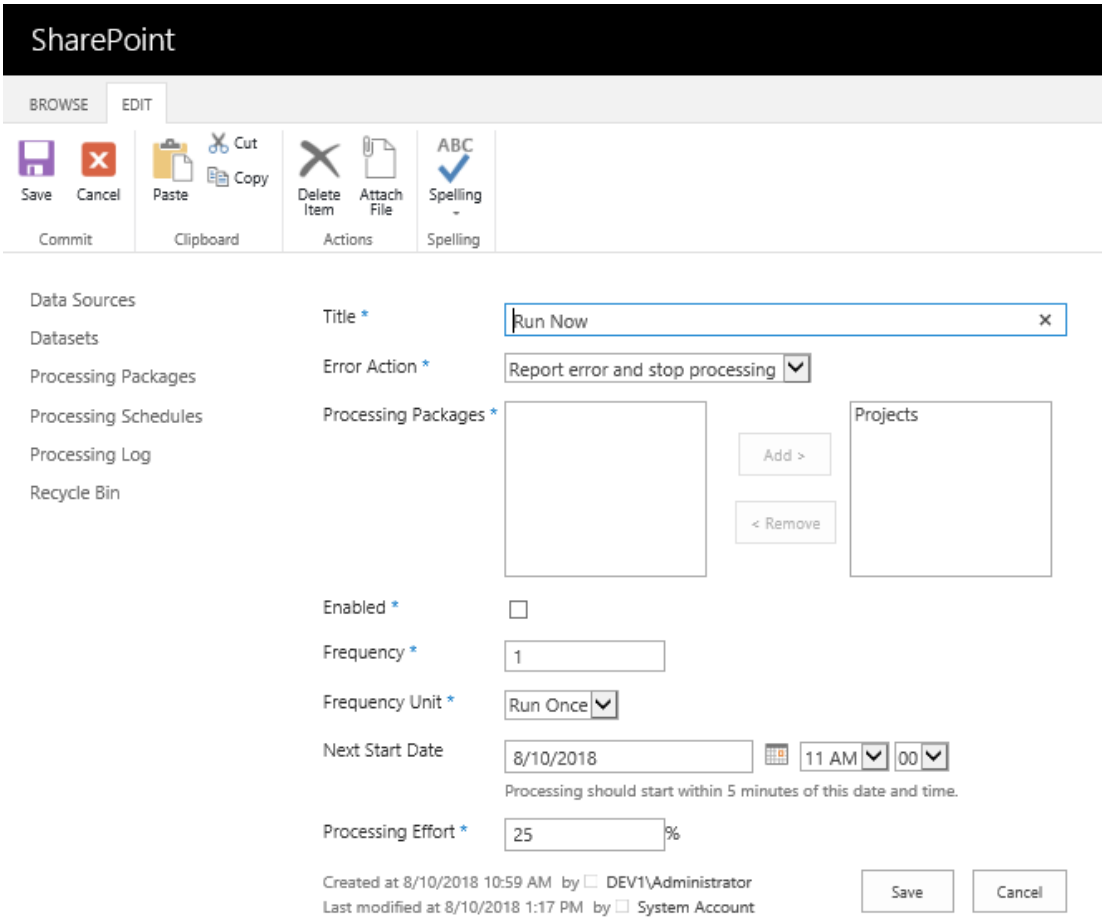

Step 1

From the quick Launch menu select "Processing Schedules"=>"Add new item".
Enter

- a value in the "Title" field
- Select the appropriate Processing Package
- Check the "Enable" box
- Enter '1' in the "Frequency" field
- Pick todays date from the "Next Start Date" field and set the time somewhere in the past.

Click "Save" to complete the Processing Schedule configuration.

Notes on Processing Schedules

The Project Integration Engine is a multi-threaded application and adjusting the Processing Effort % determines how many threads the application will use. A value of 100% will work the server as hard as it can go and is not recommended. Values between 10 and 80% seem to be the most reliable. For small processing jobs (i.e. less than 5 minutes to execute) we recommend a setting of 75%, while for long running jobs (between 10 minutes and 2 hours) we recommend tuning the effort down to around 25% to ensure other services still execute. For jobs longer than 2 hours we recommend adding an additional Application server to your SharePoint farm and shifting the Timer Jobs for PWA to the new server.

The SharePoint Timer has a minimum execution frequency of 5 minutes, so it may take this long before a job actually starts.

To execute a job immediately, open SharePoint Central Administration.
Navigate to Monitoring.
Click “Review Job Definitions”.
Select the “Project Integration Engine scheduled job processor”.
Click “Run Now”.

Processing Log

The Processing Log is a list of events reported by the Project Integration Engine when processing schedules and packages. The level of detail reported into the log is configured by the "Minimum Log Level" field on each Processing Package. The log is also cleared of the events for each Processing Package before they are executed (assuming the "Clear previous log entries" field is checked).

You can use the Processing Log to diagnose any errors or warnings reported by the Project Integration Engine, and to determine the length of time it has taken to process a package and/or schedule.

One important item to note about the Processing Log is that it contains events reported from multiple threads executing simultaneously, so if you're diagnosing events reported in it, you need to filter the list by Processing Package and then by Thread ID to highlight only the offending events. It's also a good idea to check the sort order is by ID to ensure you're looking at events in sequential order.

Monitoring a running job

There are two ways of monitoring a running job (Processing Schedule):

The first is by selecting the relevant "Processing Schedule" (using the item view, not the item edit page) and reviewing the "Status" and "Progress" fields. When a Job is completed, these fields will return to "Waiting" and "100%", but the "Next Start Date" will be incremented or the "Enabled" field will be set to "No" (Processing Schedule is disabled if it's "Frequency Unit" field is set to "Run Once").

The second is to list the running job in SharePoint Central Administration. Open SharePoint Central Administration and navigate to "Monitoring"=>"Check job status". If the job is running it will be listed as "Project Integration Engine scheduled job processor" and have a progress bar reporting the % complete.

Importing and Exporting the Configuration

Once you have configured the Project Integration Engine, you will want to migrate the entire configuration to a TEST or Production environment (because you should not be configuring this directly in Production).

Click on the "Project Integration Engine" link in the top link bar to show the home page.

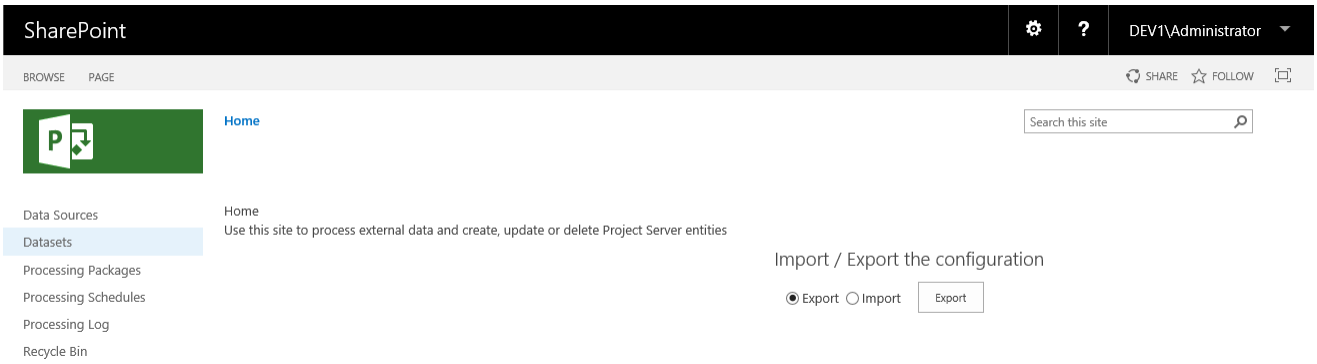

Click "Export" and save the configuration to a file.
Open your destination "Project Integration Engine" environment in a new browser window or tab.
Select the "Import" radio button.
Click "Browse...".
Select the file you just exported.
Click "Import" to complete the configuration migration.

** WARNING **

Make sure you update the Data Sources to point to the correct locations for the environment.

The entire configuration is migrated across in the process described above including the Data Sources, which means you have likely just imported DEV environment Data Sources into your Test or Production environment.

The unchecked "Overwrite existing Data Sources" check box helps to prevent this risk. Please use it with caution.

So What's Next?

Purchasing

If you've completed installing and configuring PIE 2016 and have decided to purchase a license, follow the Purchasing Guide included in your download package.

License Activation

If you have a license for PIE 2016 and wish to activate it, follow the Activation Guide included in your download package.