

CONDUCTOR

---

# User Guide

---



# Table of Contents

<b>What Is Conductor?</b>	<b>7</b>
<b>How Conductor Works</b>	<b>7</b>
<b>Conductor Entities</b>	<b>8</b>
Account	9
User	9
Project	9
Data Store	9
Agent	9
Process Group	9
Process	9
<b>Getting Started</b>	<b>10</b>
Tasks	10
<i>Sign In</i>	10
<i>Forgot Password</i>	12
<i>Sign Out</i>	14
Navigation	15
<b>Security</b>	<b>17</b>
Role-Based Security	17
<i>Account Permissions</i>	17
<i>Project Permissions</i>	17
Tasks	18
<i>Create a New User</i>	18
<i>Edit a User</i>	19
<i>Deactivate or Delete a User</i>	20
<b>Projects</b>	<b>21</b>
About Projects	21
Tasks	21
<i>Create a New Project</i>	21
<i>Assign User Permissions</i>	22
<i>Edit a Project</i>	24
<i>Delete a Project</i>	24
<b>Data Stores</b>	<b>25</b>
About Data Stores	25
<i>Types of Data Store</i>	25

<i>Cloud (Direct) vs Internal (via Agent)</i> .....	25
<i>Connection</i> .....	26
<i>Source or Destination</i> .....	26
<i>Flood Control</i> .....	26
<i>Automatic Scanning</i> .....	27
<i>Data Sovereignty</i> .....	27
<i>Data Sharing</i> .....	27
<i>Custom Data Queries</i> .....	28
<i>Filtering</i> .....	28
<i>Tagging</i> .....	28
<i>Drag and Drop</i> .....	28
<b>Tasks</b> .....	31
<i>Create a New Data Store</i> .....	31
<i>Edit a Data Store</i> .....	39
<i>Share a Data Store</i> .....	39
<i>Accept a Data Sharing Invitation</i> .....	41
<i>Pause or Resume a Sharing Agreement</i> .....	42
<i>Stop Sharing Data</i> .....	43
<i>Create a Custom Data Query</i> .....	45
<i>Edit or Delete a Custom Data Query</i> .....	46
<i>Delete a Data Store</i> .....	46
<i>Browse and Edit Data Structures</i> .....	47
<i>Manage Data Store Filters</i> .....	49
<i>Manage Data Store Tags</i> .....	50
<i>Transfer Data via Drag and Drop</i> .....	52
<b>Agents</b> .....	53
Secure Communications .....	53
Direct vs Agent Connections .....	53
System Requirements .....	54
<i>Host Computer</i> .....	54
<i>Permissions</i> .....	54
<i>Internet Access</i> .....	55
<i>Data Access</i> .....	55
<b>Tasks</b> .....	57
<i>Install an Agent</i> .....	57
<i>Configuring an Agent</i> .....	63

<i>Manually Updating an Agent</i> .....	65
<i>Uninstalling an Agent</i> .....	67
<b>Process Groups</b> .....	<b>69</b>
About Process Groups.....	69
Destination Actions.....	69
<i>Overwrite existing data</i> .....	69
<i>Append to existing data</i> .....	69
<i>Merge new and changed data</i> .....	69
<i>Merge all data (insert/update/delete)</i> .....	69
<i>Incrementally load data</i> .....	69
Tasks.....	70
<i>Create a New Process Group</i> .....	70
<i>Edit a Process Group</i> .....	70
<i>Delete a Process Group</i> .....	71
<i>Execute a Process Group</i> .....	72
<i>Schedule a Process Group</i> .....	72
<b>Processes</b> .....	<b>74</b>
About Processes.....	74
Process Options .....	74
<i>Truncation</i> .....	74
<i>Ambiguity</i> .....	75
<i>Uncertainty</i> .....	75
<i>Unusual</i> .....	75
<i>Missing Source Columns</i> .....	75
<i>Missing Destination Columns</i> .....	75
<i>Column Matching Threshold</i> .....	75
Tasks.....	76
<i>Create New Processes</i> .....	76
<i>Edit a Process</i> .....	78
<i>Delete a Process</i> .....	79
<i>Execute a Process</i> .....	79
<i>View Batch Activity</i> .....	80
<i>View Batch History</i> .....	81
<b>Column Mapping</b> .....	<b>83</b>
About Column Mapping.....	83
Smart Mapping .....	83

Defined vs Actual Column Mapping.....	85
Tasks.....	86
<i>Change column mapping</i> .....	86
<i>Prevent column mapping</i> .....	87
<b>Column Manipulation .....</b>	<b>88</b>
About Column Manipulation .....	88
Filtering .....	88
Constants .....	88
System Metadata .....	88
Expressions.....	89
Tasks.....	90
<i>Filtering a Column</i> .....	90
<i>Adding a Constant</i> .....	91
<i>Adding System Metadata</i> .....	92
<i>Adding an Expression</i> .....	93
<b>Reporting .....</b>	<b>95</b>
About Reporting.....	95
Account Groups .....	95
The Reports.....	95
<i>Aged Users Report</i> .....	95
<i>Account Summary Report</i> .....	96
<i>Data Movement Report</i> .....	97
<i>Data Security Report</i> .....	98
Using Reports.....	99
Tasks.....	99
<i>View a Report in the Web Browser</i> .....	99
<i>Download a Report in Excel, PDF or Word Format</i> .....	100
<b>Audit Log .....</b>	<b>101</b>
About Audit Logging.....	101
<b>Appendix A - Connectors .....</b>	<b>102</b>
Destination Actions.....	103
List of Connectors .....	105
<i>Application Connectors</i> .....	105
<i>File Connectors</i> .....	105
<i>NoSQL Database Connectors</i> .....	106
<i>Relational Database Connectors</i> .....	107

<i>Binary Connectors .....</i>	<i>108</i>
<b>Appendix B - System Metadata .....</b>	<b>110</b>
<b>Glossary of Terms .....</b>	<b>112</b>
<b>Contact Us .....</b>	<b>113</b>

## What Is Conductor?

Eightwire's Conductor is off-the-shelf cloud-based data sharing software.

### **Automates 80% of the work**

With traditional data integration, 100% of the effort comes from developers hand-coding. Conductor automates 80% of that effort, so skilled developers can work on the difficult technical challenges more efficiently.

### **Removes platform, location and skill silos**

No matter what platform is used, Conductor integrates data easily. It bridges the gap between the skills within an organisation and the desired goal, so difficult integration projects can be completed without years of training and experience.

### **Self-healing data feeds**

Conductor automatically manages change and fixes data errors that can cause systems to crash. For an average customer, this means several thousand errors are automatically fixed every month, saving hundreds of hours in business downtime.

## How Conductor Works

Conductor is relatively simple to use and requires only three steps to begin moving data. As a user, you will interact with Conductor entirely through the Conductor Web Portal using your internet browser:

1. Define at least one source and one destination Data Store.
2. Define one or more data transfer Processes.
3. Execute those Processes to transfer data.

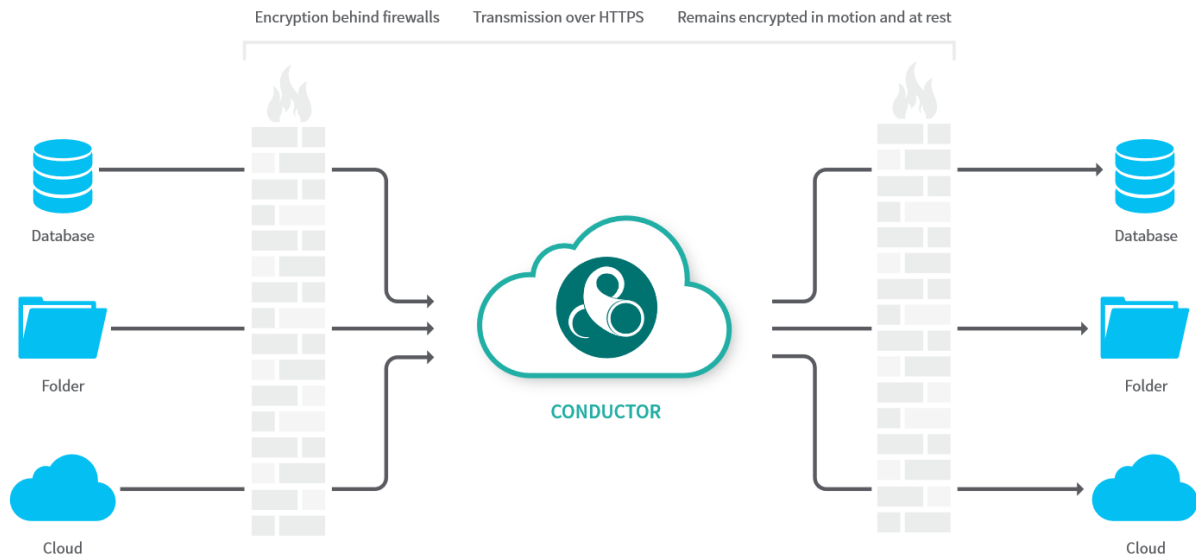
That's it.

Conductor is a cloud-based service, which means there is almost nothing to install in your own organisation. The exception to this rule is the Conductor Agent.

Most data resides on servers within corporate networks, which makes it impossible to access from a cloud-service on the internet.

Conductor uses customer-installed Agents to provide secure connections out of customer networks, allowing the Conductor cloud service to access data behind firewalls when it needs to transfer data.

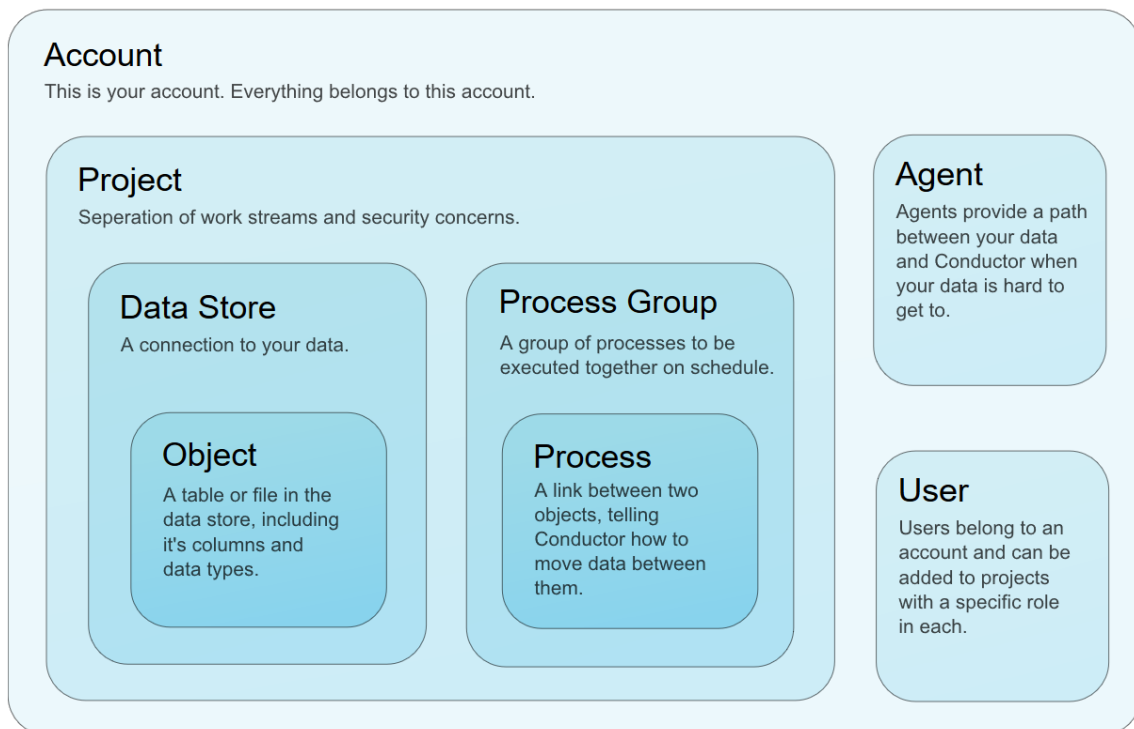
An Agent is optional, but greatly improves the security of data transfers and eliminates the need to open holes in your corporate firewall.



Conductor also supports many cloud-based data repositories, such as RedShift and Cloudant, which do not require an Agent because they are not behind a corporate firewall and can be directly accessed by Conductor's cloud-service.

## Conductor Entities

When you interact with Conductor, you will be working with different entities, such as Projects and Data Stores. It may be helpful to know how these relate to each other:





## Account

Your Account contains everything you need. Typically your organisation will have one Account. All other entities sit within your Account.

## User

An Account contains any number of Users. Typically each person accessing your Account will be a separate User. Users are assigned roles both at the Account and Project level.

## Project

Work streams within an Account are divided into one or more Projects. A project allows you to group the Users, Data Stores and Processes that relate to a particular piece of work. It is also a security boundary – only those Users with an assigned role in a given Project will have access to the entities within it.

## Data Store

The definition of a data repository. This could be a file system folder containing CSV or Excel files, a relational database containing tables, or a NoSQL database containing documents. You are providing Conductor with the information it needs to locate and connect to the data repository.

## Agent

Most data repositories are on internal customer networks, behind corporate firewalls. This makes accessing this data from cloud-services difficult or impossible. An Agent is a small piece of software that runs inside a customer's own network and provides both data access capability and a secure communications channel to Conductor's cloud services.

If a Data Store uses an Agent, the connection information given should be relative to that Agent, because it is acting on Conductor's behalf.

## Process Group

A group of related data transfer processes. Process Groups can be scheduled to run at certain times. If the objects (tables) contained within the Processes in the group relate to each other, these relationships will be respected during data transfer when the Process Group itself is executed (e.g. with relational database referential integrity constraints).

## Process

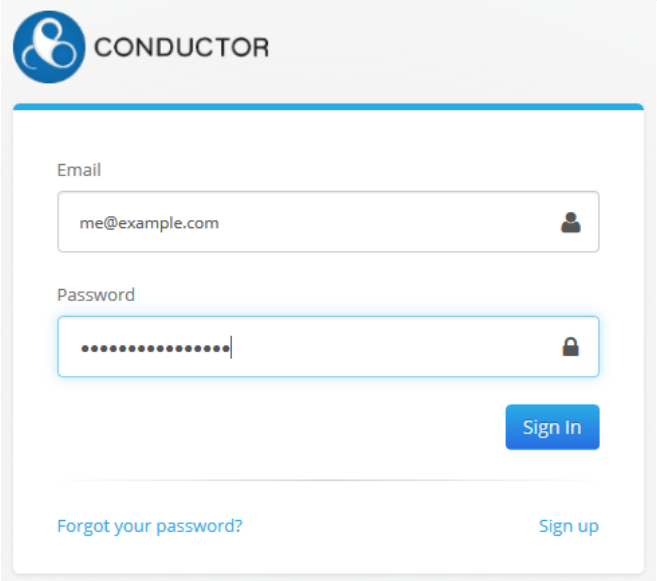
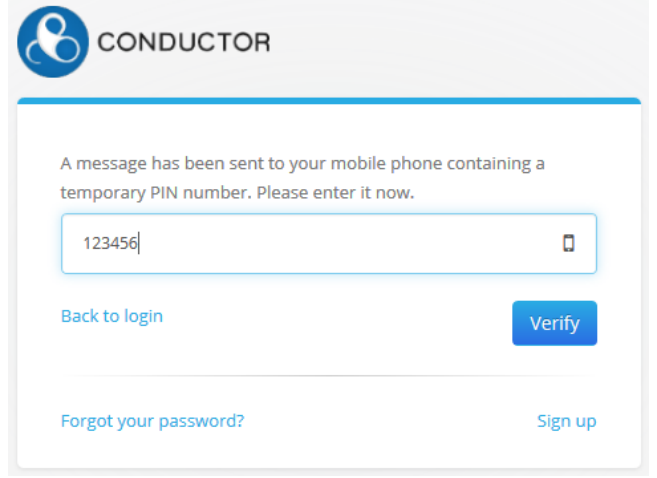
An individual data transfer definition. This defines how one object (table) is read from the source Data Store, remapped and then written to the destination Data Store.

There is no restriction on what type of source and destination Data Stores are used, provided they have been defined within the same Project. Conductor will transfer data between any of the data platforms it supports.

## Getting Started

### Tasks

#### Sign In

<p>1.</p>	<p>Open your web browser to <a href="https://conductor.eight-wire.com/">https://conductor.eight-wire.com/</a></p> <p>Enter your email address.</p> <p>Enter your password.</p> <p>If you want Conductor to remember what you typed, check the 'Remember Me' box. Do not do this on a computer that could be used by other people.</p> <p>Click 'Sign In'.</p>	
<p>2.</p>	<p>If your account requires two-factor authentication, an SMS text message will be sent to your mobile phone number, containing a single-use PIN number. Enter this PIN number on the second step of the Sign In process and click the 'Verify' button.</p>	

3. If this is the first time your account has been used, you will have to agree to the 'Terms and Conditions' before proceeding. You will only need to do this once.

If you agree, check the 'I Agree' box and click 'Continue'.

Please note: you are agreeing on behalf of your account and not just yourself. All users who use this account are bound by this agreement.

#### Terms and Conditions

### Terms and Conditions of Eight Wire Limited

By using Our Website in any way, or by purchasing any of our Services, you agree to be bound by our Terms and Conditions below.

Please note that no person under the age of 18 years may purchase our Services. We look forward to seeing you again when you are over 18.

**We are:** Eight Wire Limited  
**Our address is:** Level 1, 15 Edward Street, Wellington, New Zealand

**You are:** a visitor to Our Website / our customer

#### The terms and conditions

#### 1 Definitions

In this agreement:

"Our Website" means eight-wire.com (including www.eight-wire.com)

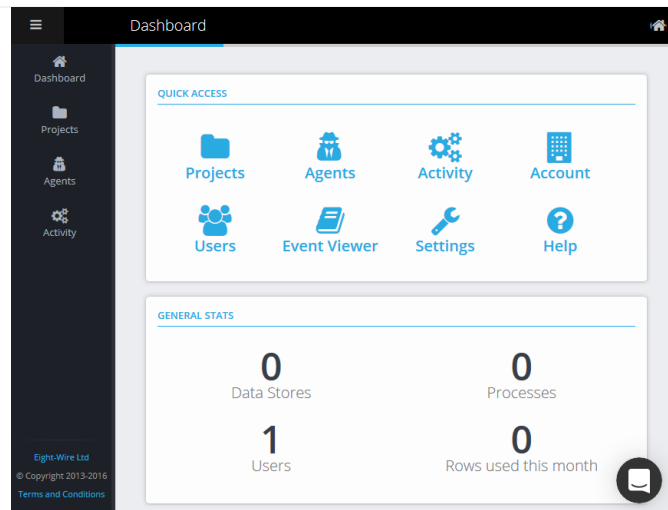
☒ I Agree to the Terms and Conditions.
 

[Back](#)
[Continue](#)

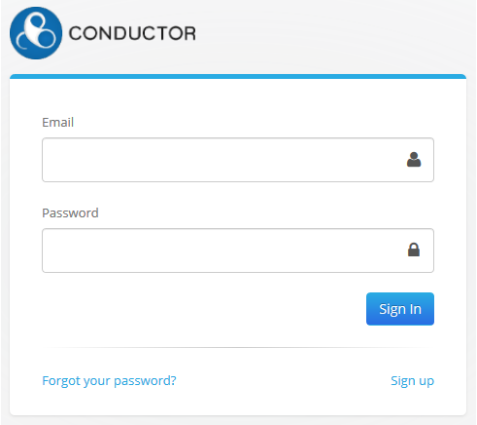
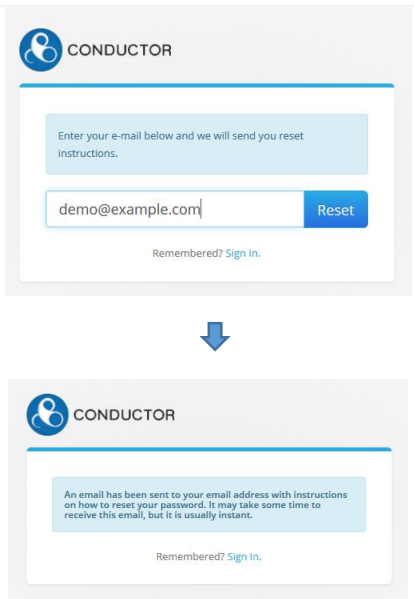
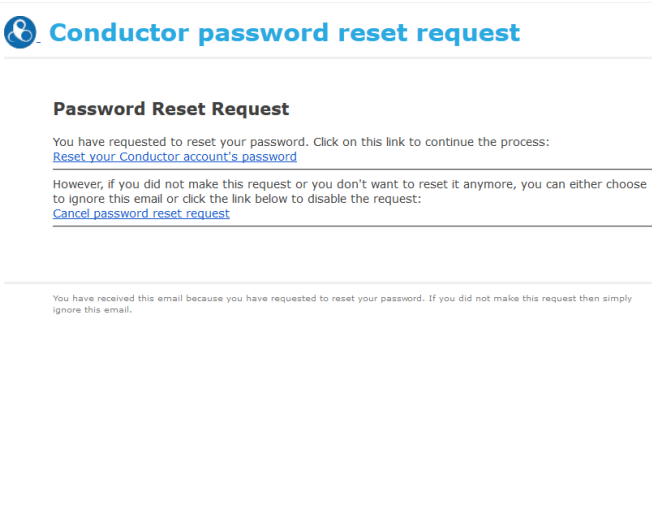
4. You should arrive at the Dashboard as shown.

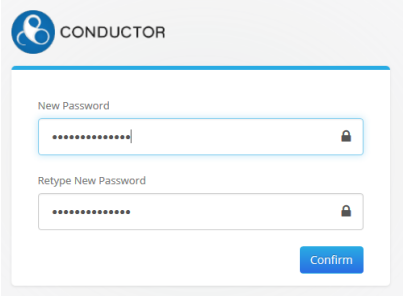
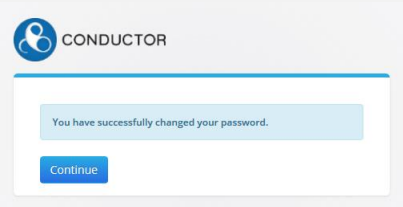
Congratulations, you are now signed in and can begin using Conductor.

If you have forgotten your password, use the 'Forgot your password?' link on the Sign In page.



## Forgot Password

1.	<p>From the main login screen, click 'Forgot your password?'</p>	
2.	<p>Enter your email address. Click 'Reset'.</p> <p><i>You will see a confirmation that an email will be sent, provided the email address entered was valid.</i></p>	
3.	<p>Check your email. You should have received an email entitled 'Conductor Password Reset'.</p> <p>Open the email and click 'Reset your Conductor account's password'.</p> <p><i>If you have changed your mind you can cancel the password change by clicking the 'Cancel password reset request' link in the email instead.</i></p>	

4.	Enter your new password twice and click 'Confirm'.	
5.	<p>You will see a confirmation that your password has been changed.</p> <p><i>Click 'OK' – you will be taken to the login screen, where you will be able to login with your email address and new password.</i></p>	

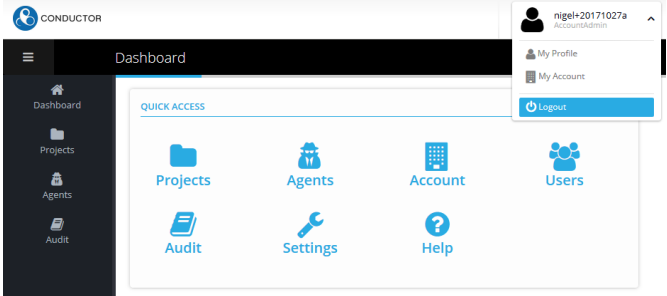
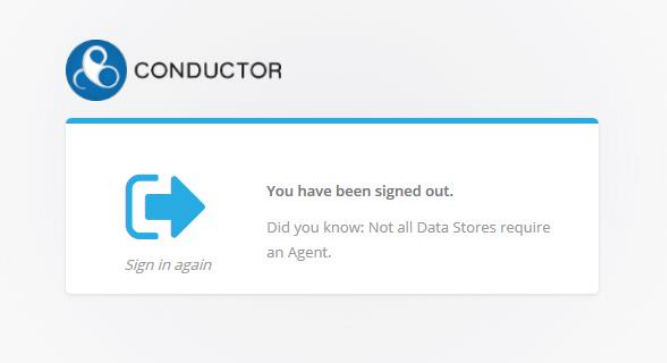
### Password Complexity Requirements

All Conductor passwords must adhere to the following rules:

- Must be either 16 characters or more long, or must have at least 10 characters including three of uppercase, lowercase, numeral and symbol.
- Must not be the same as any of the previous 12 passwords used.

Your password needs to be changed at least every 90 days.

## Sign Out

1.	<p>At any time you can sign out of Conductor by clicking 'Sign Out' from the menu in the top toolbar.</p> <p>Alternatively, at any time, if you do nothing for 10 minutes, you will be automatically signed out.</p>	 <p>The screenshot shows the Conductor dashboard interface. In the top right corner, there is a user profile dropdown menu. The menu includes options for 'My Profile', 'My Account', and a prominent 'Logout' button.</p>
2.	<p>You are now signed out of Conductor.</p> <p><i>To sign in again, please go to the 'Sign In' section of this manual.</i></p>	 <p>The screenshot shows the Conductor sign-in page. It features the Conductor logo at the top. Below the logo, there is a large blue arrow pointing right, with the text 'Sign in again' underneath it. To the right of the arrow, the text reads: 'You have been signed out. Did you know: Not all Data Stores require an Agent.'</p>

**When you sign out, any unsaved changes will be lost.**

## Navigation

**This is the Dashboard.** It is your ‘home’ screen. From here you can get to any area of Conductor. Depending on your account role, you may not see all of the navigation buttons shown here:

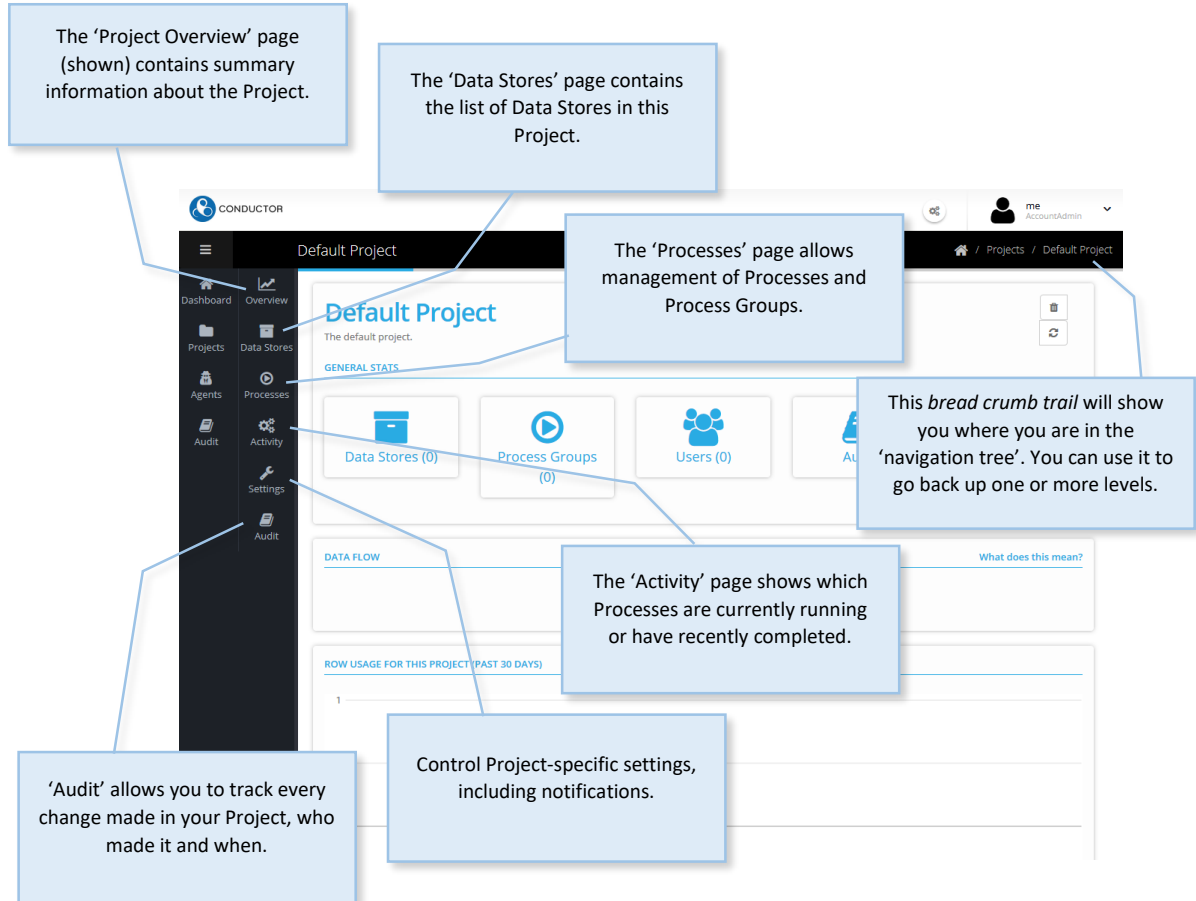
The screenshot shows the Eightwire Conductor Dashboard. A dark sidebar on the left contains navigation icons for Dashboard, Projects, Agents, and Audit. The main content area has a 'QUICK ACCESS' section with icons for Projects, Agents, Account, Users, Audit, Settings, and Help. To the right is a 'GENERAL STATS' section showing 0 Data Stores, 0 Processes, 1 Users, and 0 Rows used this month. Below this is a 'ROW USAGE FOR ALL PROJECTS (PAST 30 DAYS)' chart. Callout boxes provide the following information:

- Top Left:** To get back to the dashboard from anywhere, click the ‘Dashboard’ icon or the Conductor logo.
- Top Center:** ‘Projects’ will take you into the Projects area, where you will likely spend most of your time in Conductor.
- Top Right:** This menu allows access to user and account settings, and **Sign Out**.
- Bottom Left:** The ‘Agents’ section is where you will find information about your Agents, such as online / offline status.
- Bottom Center:** ‘Audit’ allows you to track every change made in your account, who made it and when.
- Bottom Right:** The Dashboard shows general information about your account.

The left-hand navigation here is shown as a single column of icons. When you are not working within a Project, you will see this navigation panel. When you click into a specific project, you will see a more detailed navigation panel.

**When working within a Project**, you will see an extended navigation panel on the left-hand side of the screen.

The main window content will change as you navigate through Conductor, but you can always get back to one of the main Project areas, or main Account areas using the left-hand navigation panel.





## Security

There are many aspects to security in Conductor. In this manual, we will only cover user permissions and roles in any detail. If you would like more information about other aspects of security, such as data security, data retention, data sovereignty or Agent security, please get in touch with Eightwire support staff, who will be happy to help.

### Role-Based Security

Each Conductor Account contains one or more Users and one or more Projects. Conductor uses *role-based security* to determine what users can see and do within the Account and within each individual Project.

#### Account Permissions

Each user is allocated one role at the Account level:

- **Account Admin**  
*Full access to all Account-level information. Can create Users and Projects, assign users to Projects and manage Agents.*
- **Account Viewer**  
*Read-only access to most Account-level information.*
- **Account User**  
*Limited read-only access to Account-level information. Usually operates at a Project level only. Most users in a larger organisation will have this Account Role.*

**You should assign each user the lowest level of access necessary. This is especially true at the Account level as these permissions are far-reaching and could have a big impact on your organisation.**

You must have at least one user with the Account Admin role.

#### Project Permissions

Most users in your organisation will be working on specific projects. Each user is assigned to various projects by either an Account Admin or Project Admin. What the user can see and do in each Project is governed by their project role:

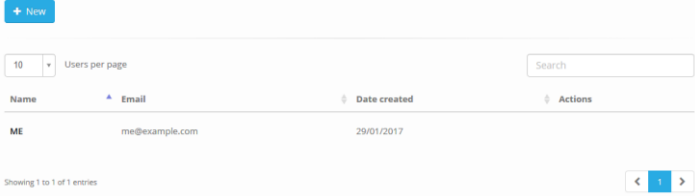
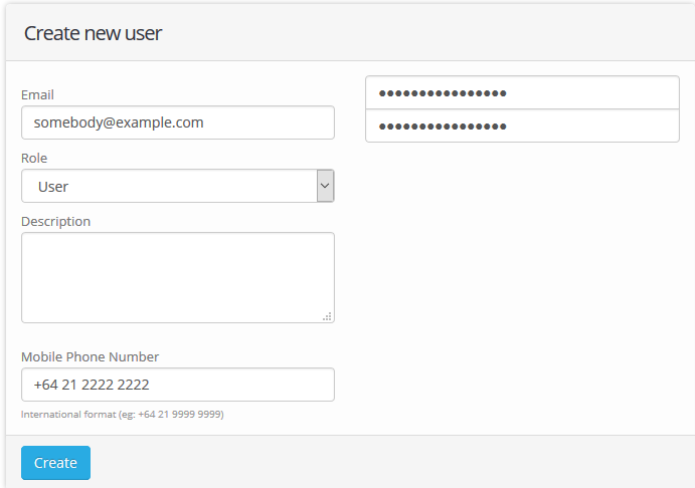
- **Project Admin**  
*Full access to all Project-specific information. Can assign users to the Project, create Data Stores, Agents and Processes.*
- **Project Author**  
*Full access to Processes and Process Groups only. Can create, schedule and execute Processes.*
- **Project Viewer**  
*Read-only access to the Project Overview only.*

In essence, an Admin defines where the data is coming from and where it is going to (Data Stores), the Author defines how the source and destination is mapped together and schedules when data is moved, and Viewers oversee the Project at a high level but have no access to anything sensitive.

## Tasks


To create, edit or delete users, you need to be on the **Users** page. From the Dashboard, either click on the 'Users' Quick Access icon, or from the menu at the top-right, click 'My Account', then on the next page click the 'Users' button at the bottom right.

### Create a New User

1.	Click the 'New' button:	 <p>The screenshot shows the 'Users' page interface. At the top left is a blue '+ New' button. Below it is a table with columns: Name, Email, Date created, and Actions. The table contains one entry with Name 'ME', Email 'me@example.com', and Date created '29/01/2017'. At the bottom right, there is a pagination control showing '1' of 1 entries.</p>
2.	<p>Enter an email address.</p> <p>Select an <b>Account Role</b> (<i>see the Security section in this manual</i>).</p> <p>Enter a mobile phone number (must be in international format).</p> <p>Enter a password and retype the same password again.</p> <p>Click 'Create'.</p>	 <p>The screenshot shows the 'Create new user' form. It includes fields for Email (with example 'somebody@example.com'), Role (a dropdown menu set to 'User'), Description (a text area), and Mobile Phone Number (with example '+64 21 2222 2222'). There are two password fields represented by dots. A 'Create' button is at the bottom.</p>
3.	An email will be sent to this new user, asking them to verify their email address and sign in. <i>For more information on signing in, see the Sign In section in this manual.</i>	

## Edit a User

1. Hover your mouse cursor over the User you wish to edit. Notice the buttons that appear for that user. Click on the 'Edit' button that appears next to the User.



2. Modify the details you wish to change and click 'Save'.

The screenshot shows the 'Edit personal details' form. It includes the following fields:
 

- Email: someone@example.com
- Role: User (dropdown menu)
- Description: A large text area.
- Mobile Phone Number: A field with a '+' icon and a note 'International format (eg: +64 21 9999 9999)'.
- Change password section:
  - Current password: A text field.
  - To change your password, enter your current one then your new ones below.
  - New Password: A text field.
  - Retype new password: A text field.

 At the bottom is a blue 'Save' button.

### Password

If you do not wish to change the password, leave these fields blank. Or, enter the current password and then enter the new password (twice) to change it, then click 'Save'.

### Mobile Phone Number

If your account requires two-factor authentication, you will be required to enter a valid mobile phone number, in international format. That means the number should start with '+' followed by the country code (e.g. 64 for New Zealand), followed by the area code and phone number without leading zeros. For example, a New Zealand mobile phone number might look like this:

+64 21 987 6543.

You will be sent an SMS text message to confirm that the mobile phone number you provided is correct.

### Email

If you change your email address, you will be sent a confirmation email. You will need to click on the link in the email to confirm that this is your correct email address.



## Deactivate or Delete a User

3. Hover your mouse cursor over the User you wish to deactivate or delete. Notice the buttons that appear for that user.

Click on the 'Deactivate' button that appears next to the User to keep the user but deactivate them, preventing anyone signing in with that user account. To re-activate them, click the 'Activate' button.

Click on the 'Delete' button and confirm when asked to permanently delete the user. **This cannot be undone.**

*Deactivating or Deleting a User does not remove their audit history, so even after a user account is turned off or removed entirely, you will still be able to access a complete history of their actions in Conductor.*



# Projects

## About Projects

A Project is a workspace with assigned Users, Data Stores and Processes. It is an easy way to segregate different work streams within your Account.

Users can have a different role for each project they are assigned to, or not be assigned to any Projects at all. This provides control over who will work on which projects and in what capacity. Account Admins have permission to create or delete Projects and assign Users to Projects, including themselves. A user cannot participate in a Project if they are not explicitly assigned to it, including Account Admins.

Data Stores are unique to each Project. If you wish to make the same Data Store available to multiple Projects, you will need to either duplicate it, or share it.

The only things not unique to each Project are the Agents. These belong to the Account, but are used by the Data Stores within a Project.

When your Account is first created, an empty Project named 'Default Project' will automatically be created for you. Feel free to rename it to anything you like and start using it, or simply delete it and create your own.

## Tasks

To create, edit or delete Projects, you need to be on the **Projects** page. From the Dashboard, click on the 'Projects' icon on the left-hand navigation bar.

### Create a New Project

1. Click the 'New' button to create a new Project.

#### Projects



**Default Project**

The default project.

+ New

2. Add users to the project with specific roles.

*For an explanation of the roles, see the Security section in this manual.*

If you wish to participate in this project, you will need to add yourself to this project.

Click 'Next'.

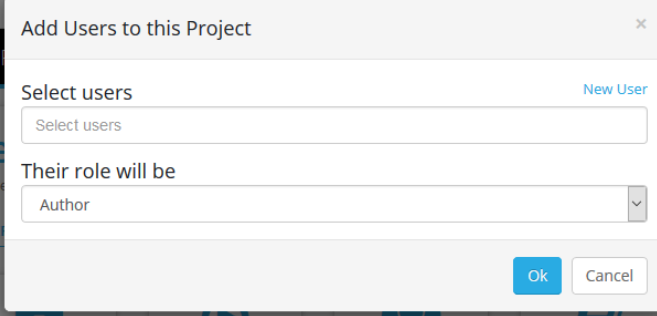
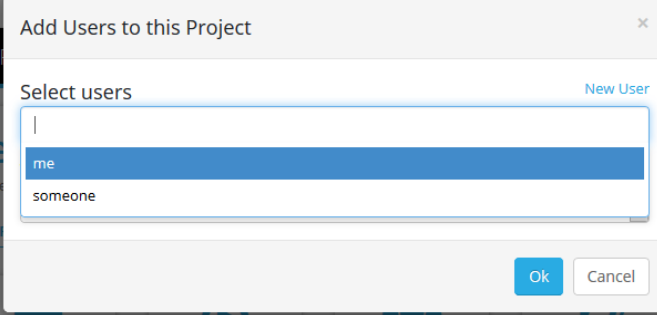
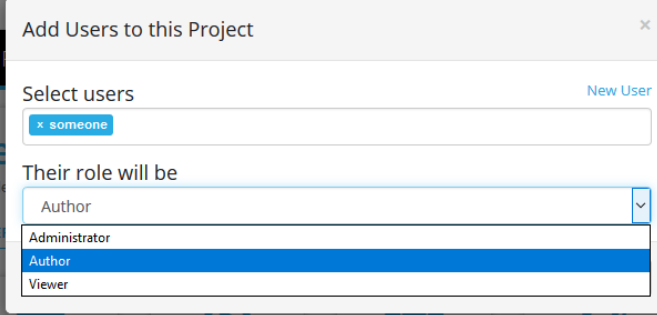
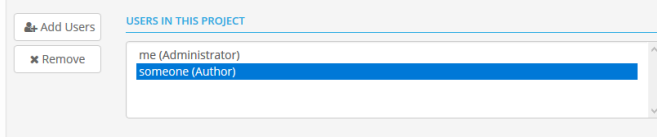
3. Give your Project a name and a description.

Click 'Create'.

## Assign User Permissions


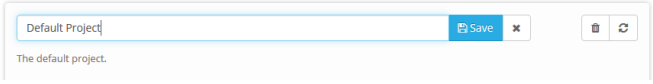


1. From the *Project Overview* page for your existing Project, click the 'Users' icon.

This will show the 'USERS IN THIS PROJECT' panel, where you can manage user access to this project.

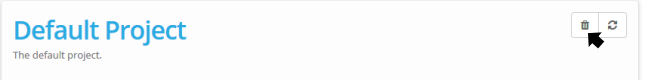
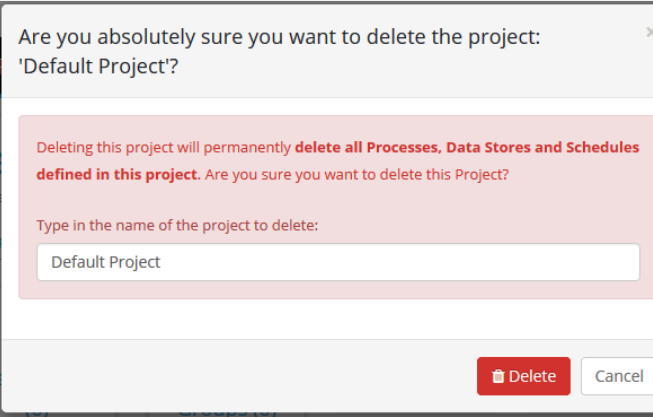
2.	<p>To add a user, click the 'Add Users' button. A popup window will appear, allowing you to select users and a role to assign to them.</p>	
3.	<p>In the 'Select users' box, start typing the user name you are looking for and a list of matches will appear. Choose the correct username. If you wish, you can add more than one user to be assigned this role.</p>	
4.	<p>Click the 'Role' box. A list of available roles is shown. Select the role you wish to give the user(s) selected.</p> <p><i>For an explanation of the roles, see the Security section in this manual.</i></p> <p>Click 'OK' to assign these user roles to the Project.</p>	
5.	<p>To remove a user role, select the user role you wish to remove and click 'Remove'.</p> <p>Click 'Yes' to confirm, when prompted.</p>	

## Edit a Project

To rename a Project or edit its description, follow these steps. To modify the contents of the Project, see the *Data Store* and *Process* sections of this manual.

1.	To rename a project, hover you mouse cursor over the Project's name. This will show a small pencil icon – click this icon to rename the Project. When done, click 'Save'.	 
2.	To change the Project's description, hover over the subtext below the title. This will show a small pencil icon – click this icon to change the Project's description. When done, click 'Save'.	 

## Delete a Project

1.	To delete a Project, click the 'trash' icon.	
2.	<p>When prompted, type the complete name of the Project into the warning dialog box that shows.</p> <p><i>This added step of typing the name is there to make sure you do not accidentally delete a Project you did not intend to delete.</i></p> <p>Click 'Delete'.</p>	
	<p><b>Deleting a Project is permanent.</b></p>	



## Data Stores

### About Data Stores

A Data Store in Conductor is a definition of how Conductor should locate and connect to a data repository on your behalf. Data Stores are intended to provide a level of abstraction away from the platform-specifics of data access and data types. Generally, once a Data Store has been configured, you can treat all Data Stores the same way – they contain data. Nothing more. Nothing less.

### Types of Data Store

Conductor supports a wide range of common data repositories (connectors), including:

- Access
- ANSI SQL
- DB2 (Generic)
- DB2 (iSeries)
- Cloudant
- CouchDB
- Delimited Text (*CSV and TSV files*)
- DocumentDB
- Excel (*XLSX files*)
- Fixed Width (*Padded files*)
- Hadoop (Hive)
- Interbase
- MongoDB
- MYOB
- MySQL
- Netezza
- Oracle
- PostgreSQL
- Redshift
- Salesforce (*API*)
- SQL Server
- Sybase

Different types of data repository will require different types of connection information. Aside from the connection details being different, you should treat all Data Stores the same way – they are generic data repositories.

### Cloud (Direct) vs Internal (via Agent)

A data repository can reside in the cloud or on a corporate network, behind a firewall.

If Conductor's cloud service can make a direct connection to a cloud repository, then no Agent is required.

If a data repository is behind a firewall, the customer can install an Agent on the same internal network. The Agent acts like a data proxy allowing Conductor to connect to the data repository via the Agent.

One Agent can act as a proxy for multiple data repositories, providing they are all accessible from the same Agent.

## Connection

This information tells Conductor where to find the data repository and how to connect to it. The connection information is also relative to where the connection is being made from – a direct connection will require a fully qualified domain name and to be accessible over the internet, while an Agent connection will require the connection details to be relative to the Agent, not the cloud.

For example, a direct database connection to SQL Azure (with the SQL Server connector) may require a connection string similar to:

```
Server=tcp:myserver.database.windows.net;Database=myDataBase;  
User ID=[LoginForDb]@[serverName]; Password=myPassword;  
Trusted_Connection=False; Encrypt=True;
```

While an Agent-connected CSV folder may require a file path similar to:

```
C:\Data\MyCSVFiles\
```

The second example is relative to an Agent, and points to a local folder on the same computer with the Agent.

Not all connectors support both direct and Agent connections. Most cloud-based repository types only support direct connections. File-based connectors generally only support Agent connections. Other connectors will support both types of connection.

## Source or Destination

A Data Store can act as either a source or a destination, but not both. When creating the Data Store, you will need to make a conscious decision about whether you intend to read or write data to this Data Store, which reduces the likelihood of accidentally overwriting your source data.

## Flood Control

Conductor's cloud service is designed to service multiple customers simultaneously and as such its servers have considerable combined distributed processing power. If all of this power were brought to bear on one database without some restrictions, it might be possible for Conductor to unknowingly overload your database server while it tries to read data from 100 tables simultaneously, sometimes called 'flooding'.

Each Data Store allows you to configure the maximum number of simultaneous read or write options that can be performed against that Data Store's underlying repository.

## Automatic Scanning

When you create a Data Store, one of the first things that Conductor does is to look into the underlying repository and gather information on its tables, columns and data types. This is called *Discovery and Scanning*. When this is complete, Conductor has a complete cached copy of this information, which assists with smart mapping and detecting unexpected data structure changes.

You can choose whether to periodically allow Conductor to rescan and refresh its cache or only do this when requested.

If your underlying repository data structures change, you should rescan the Data Store to ensure you are working with accurate data structures.

Conductor will automatically scan each table during a data transfer and report a warning if the cached data structure deviates from the actual data structure in a way that is incompatible.

## Data Sovereignty

Conductor contains a number of *data processing regions*. Currently, there is one data processing region in New Zealand and one in the United Kingdom.

A Data Store can optionally be locked to one of these regions. For example, a data transfer Process using a New Zealand region-locked Data Store will only send data to or receive data from one of Eightwire's New Zealand processing servers. If both source and destinations are region-locked to different regions, the data transfer cannot proceed.

If a region is not selected for either source or destination Data Store, then all data transfer Processes between them will be allocated to any Eightwire processing server regardless of location.

## Data Sharing

Conductor can be used to transfer data between your own Data Stores, but sometimes it may be desirable to share data with another organisation. To do this, each Data Store can be shared with other Conductor Accounts.

If an external Data Store is shared with you, you will be able to transfer data either to or from it, but you will not be able to configure it in any way. The owner of the Data Store retains complete control and can terminate the sharing agreement at any time.

Transferring data between shared Data Stores is a very secure way to share data between organisations.

A sharing agreement can be paused (put on hold) or deleted entirely, to stop current and future data transfers from taking place via that sharing agreement. Pausing simply disables the agreement but does not remove it, and it can be resumed again later. Deleting a data sharing agreement is permanent and will impact all data transfer Processes build on that sharing agreement.

**IMPORTANT!** Remember that receiving parties will only be allowed to see what you allow them to see. Please make sure that for databases you make use of the filtering functionality in order to lock down the tables, fields and rows that you allow Conductor to access for each project. Alternatively, you can extract data into a file or use a spreadsheet which should reside in a separate folder for each Data Store.

## Custom Data Queries

For Data Stores whose underlying repository is a relational database, such as SQL Server or Oracle, you can optionally add SQL queries to your Data Store.

These SQL queries work like any other source database table and you can use them in any data transfer as if they were a table or view. Queries are executed and the results treated as if they were data from a table like any other. In this way, complex queries or stored procedures can be used as source data.

## Filtering

If you don't want all of the data in a table to be transferred or available for transfer, you can add row-level filters to any table in a Data Store, regardless of the type of repository used.

Filters work at both the source and destination as well. A source table filter will restrict the data read from the source table, then the destination table filter will further reduce the data that is written to the destination table. A third filter can also be applied to each Process itself.

You can also disable whole columns if you don't want them to be available for use.

One big advantage of filtering is to control what data is made available to the external organisation you have shared your source Data Store with.

Filtering is applied equally to any type of Data Store, so you do not need to be concerned with how to pre-filter the underlying repository before Conductor access it.

## Tagging

Data structures can be tagged with arbitrary labels, giving them additional meaning. A list of tags is curated by the Project owner, or may be subscribed from an external Account.

Tags can be used as a cross-repository data dictionary, disambiguating the purpose of data structures.

Each tag can optionally be flagged as *notifiable*. Objects tagged with a notifiable tag can be monitored and tracked more closely.

Tags have relationships with each other, allowing tags to be inferred based on the relationships between tags without all related tags needing to be added to a structure. This is especially useful when determining which data structures are notifiable.

The 'Personally Identifying' tag is built-in and is always flagged as notifiable.

For example, we create a tag called 'Date of Birth' and relate it to the built-in Personally Identifying tag. We then apply the new tag to a 'DOB' column in a data structure. This would result in the DOB column being tagged with the 'Date of Birth' disambiguation tag and by association it would be notifiable as well.

## Drag and Drop

Also called 'drop file', these Data Stores are designed to make uploading CSV or Excel data from one organisation to another as easy as possible. Drag and Drop Data Stores can only be used as a source

(not a destination). These Processes are executed by literally dragging and dropping a file onto a *Drag and Drop area* on the Project Dashboard. Drag and Drop Processes cannot be executed via schedule, by pressing the 'execute' button or via an API call.

Drag and Drop does not require any Agent installation for the source Data Store.

A typical use-case for Drag and Drop would be when a central organisation or agency receives data from one or more small organisations or NGOs. Where data resides in CSV or Excel format, those small organisations or NGOs simply need to be given access to a shared Drag and Drop Data Store and they can then upload data from their data files to the central organisation or agency. This is often called a 'hub and spoke' model.

The procedure for setting up a data transfer between two organisations when using Drag and Drop is slightly different from how you would setup other shared data transfers. Broadly, you should follow these steps:

Sending Organisation	Receiving Organisation
<p>Accept the sharing invitation and link the shared Data Store to an existing Project.</p> <p>Locate the <i>Drag and Drop area</i> on your Project dashboard for this shared Data Store. Drag an appropriate file onto this area to initiate a data transfer.</p>	<p>Create a destination data store (any type). Set Data Sovereignty, if required.</p> <p>Create a Drag and Drop source Data Store and reference the destination Data Store when asked.</p> <p>Build one or more Processes to transfer data between your source and destination Data Stores.</p> <p>Share your source Drag and Drop Data Store with the <i>Sending Organisation</i>.</p>
	<p>The data transfer runs in this account. The log of its success or failure is here – the Sending Organisation has no access to this log.</p>

The file must be of the correct type for the Drag and Drop Data Store – for example, if you use a *Delimited Text Dropfile* type, the Sending Organisation must drop CSV files onto it. If you use an *Excel Spreadsheet Dropfile* type, it must only be used with Excel Spreadsheets (.xlsx files).

The Sending Organisation is also only able to drop files that exactly or closely match the structure of a Process that has already been created to transfer data. When a file is dropped, Conductor

evaluates the structure of the file and chooses the best existing Process to use for the data transfer. If no matching Process is found, the file is rejected. For this reason, it is suggested that the Receiving Organisation provides guidance to the Sending Organisation on the columns and data types that are expected for each file being transferred.

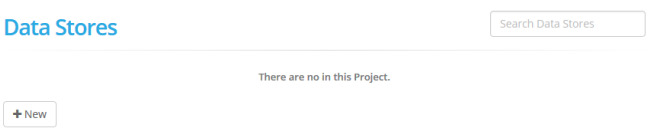
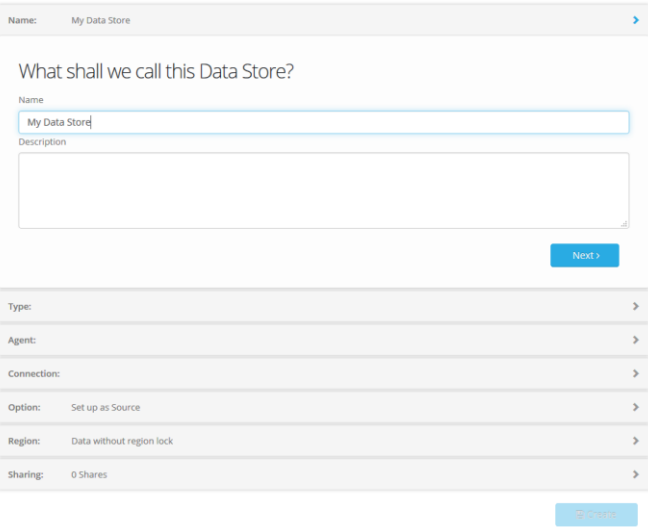
It is important that the structure of each table used in a Drag and Drop Data Store is different from the other tables used – Conductor uses this structure to determine which Process should be used to transfer the data from the file being uploaded. If multiple tables have similar or identical structures, more than one Drag and Drop Data Store should be created, linked and shared, giving the user a choice of which Drag and Drop area to drop a file onto for clarity.

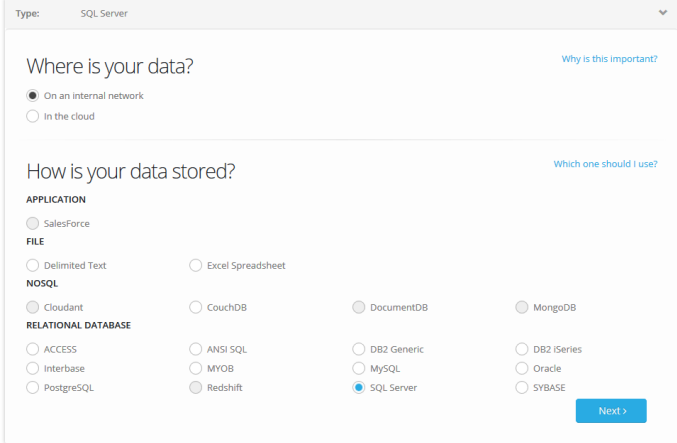
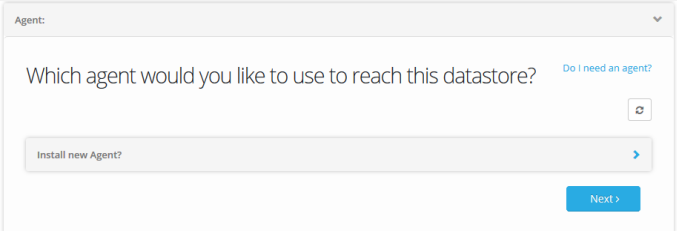
## Tasks

To create, edit or delete Data Stores, you need to be on the **Projects Overview** page or any other page belonging to the Project. From the Dashboard, navigate to *Projects*, choose one Project from the list, and then navigate to *Data Stores* – either from the left-hand navigation toolbar or by clicking the Data Stores icon on the Project Overview page.

### Create a New Data Store

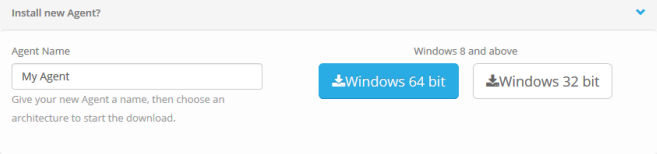
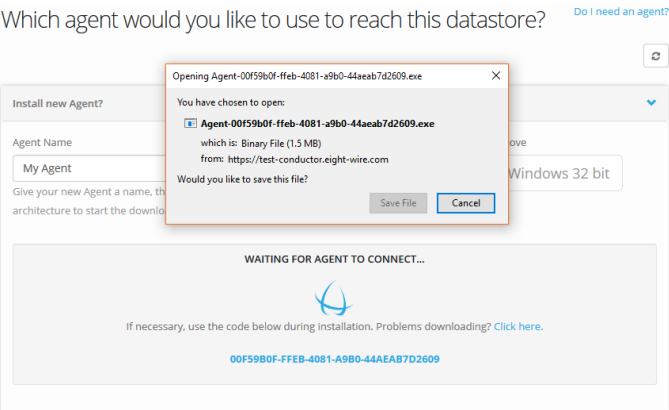
For an explanation of any of the terms or settings mentioned here, see the previous pages in this section of the User Manual.

1.	From the Data Stores list page, click the 'New' button.	
2.	<p>The New Data Store wizard has 7 panels that you will need to click through, providing information as necessary.</p> <p>First, give the Data Store a name, and optionally a description.</p> <p>Click 'Next' to move to the next page.</p> <p><i>If you share this Data Store externally, this is the name that will be shown to the external party, so it's a good idea to make it descriptive.</i></p>	

3.	<p>Choose where and how your data is stored.</p> <p>Choose whether your data resides in the cloud or on an internal network (such as behind a corporate firewall).</p> <p>Choose how your data is stored – what type of format it's in or what type of database system it's in.</p> <p>Click 'Next' to move to the next page.</p>	
4.	<p>If your data resides on an internal network, you will need to route this Data Store through an Agent, installed on the same internal network.</p> <p>If you already have an existing Agent that you would like to use, select it here, otherwise create and download a new Agent (shown here).</p> <p>Click the 'Install new Agent' header to expand this sub-panel.</p>	

**You must have *local administrator* privileges on this computer to be able to install the Agent.**




<p>5.</p>	<p>Enter a name for your new Agent.</p> <p>Click the download button for the version you want to download.</p> <p><b><i>Most users should download the 64 bit version unless using an older operating system. If you are not sure, contact the Eightwire support desk.</i></b></p>	
<p>6.</p>	<p>Your browser should either immediately download the Agent or prompt you to download it – <i>this will vary with different web browsers.</i></p> <p>Proceed with the download.</p> <p>This web page will wait for confirmation of a successful connection between the Agent and the Eightwire cloud service. You will need to complete the installation processes first.</p> <p><i>The download filename may contain a unique identifier which looks like random letters and numbers – this is normal.</i></p>	

7. Locate the downloaded Agent installation file and execute it.

Again, the procedure for doing this will be different from browser to browser. *These screenshots show the 'Downloads' window in Firefox, as an example.*

**When the Agent Installer starts, follow the prompts to complete the installation process and start the Agent service.**

*For more information on how to install an Agent, see the **Agent** section in this manual.*



8. When the Agent has been successfully installed and has successfully connected to the Eightwire cloud service, the Data Store creation web page will show a confirmation.

Make sure the newly installed agent is selected under 'Online Agents'.

Click 'Next' to move to the next page.

Which agent would you like to use to reach this datastore?

Do I need an agent?

Online Agents:

☒ My Agent

Install new Agent?

Agent Name

Enter a new Agent name

Give your new Agent a name, then choose an architecture to start the download.

Windows 8 and above

☒ Windows 64 bit

☐ Windows 32 bit

☒ AGENT SUCCESSFULLY INSTALLED.

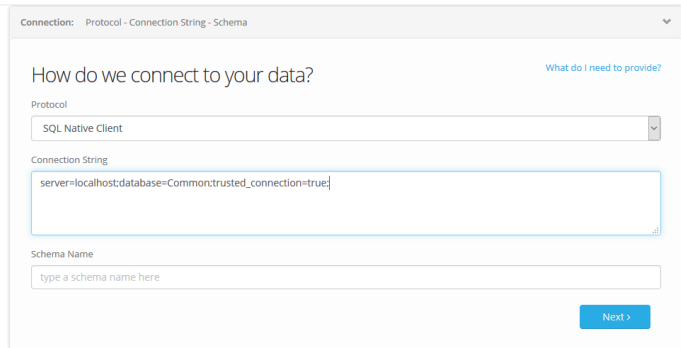
Next >

9. The *Connection* panel will change, depending on what type of Data Store you have selected - the fields shown are relevant to that type of Data Store.

*In the example shown, we are asked to provide a database protocol and connection string because we selected SQL Server as our Data Store Type for this example.*

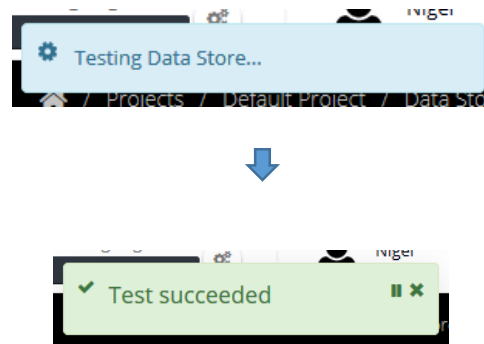
Enter the required information.

Click 'Next' to move to the next page.



10. Eightwire now has enough information to test the connection from Eightwire's cloud service through to your database. You will see a message popup to this effect. The message should change to 'Succeeded' within a few seconds. If it does not or there is a problem, please click on the 'Connection' header, re-check the connection details and click 'Next' again.

When you have successfully established a connection, click 'Next' to move to the next page.



11.

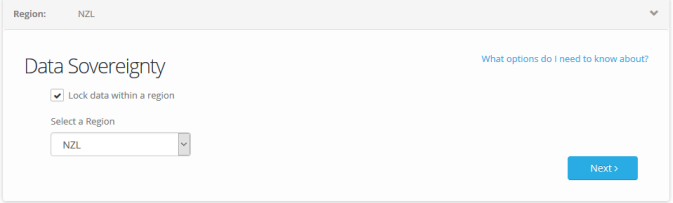
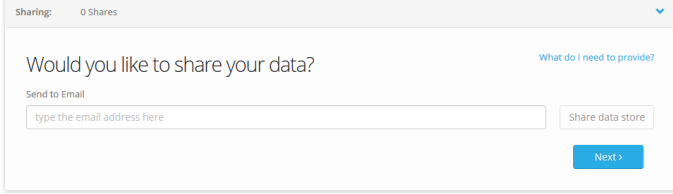
Choose whether this Data Store will be a source or a destination (i.e. you will be reading from it or writing to it).

You may also wish to adjust the 'flood control' slider to restrict the number of tables being read or written to concurrently. If your server or database has performance issues, you may wish to reduce this slider's position to reduce the load Eightwire will place on it.

If you would like Eightwire to automatically keep your Eightwire Processes in sync with the tables in your database (recommended), select 'Yes' on the bottom option. If you select 'No', you will need to regularly tell Eightwire to re-scan your Data Store to make sure it stays current with any changes to tables and columns in your database.

Click 'Next' to move to the next page.

The screenshot shows a configuration window titled "Option: Set up as Source". The main heading is "How would you like to configure this datastore?". Below this, under "This Data Store:", there are two radio buttons: "can be a source" (which is selected) and "can be a destination". Below these is a slider for "Maximum concurrent reads: 5", with a blue bar and a dot indicating the current position. To the right of the heading is a link: "What options do I need to know about?". The second section is "Would you like to keep this datastore updated?". Below this, it says "Scan process will keep the datastore's metadata updated:" followed by two radio buttons: "Yes" (selected) and "No (manually updated by user)". A blue "Next >" button is located at the bottom right of the window.

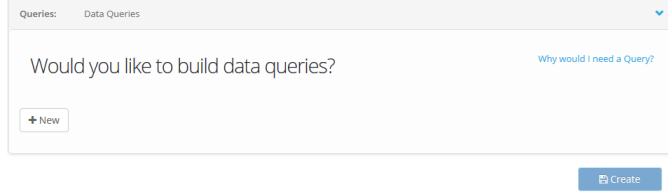
<p>12.</p>	<p>If you require your data to remain within one of Eightwire's data sovereignty regions, check the box and select the region here.</p> <p><i>Note: You may be legally obliged to use a data sovereignty region if you are a government entity or are sharing data with a government entity.</i></p> <p>Click 'Next' to move to the next page.</p>	
<p>13.</p>	<p>To setup data sharing for this Data Store, you may do this here.</p> <p><i>Remember to <b>save</b> the Data Store after sharing it otherwise the invitee will not be able to share it.</i></p> <p><i>For more information on sharing data stores, see the <b>Data Sharing</b> section in this manual.</i></p> <p>Click 'Next' to move to the next page.</p>	

14.

To create Custom Data Queries, you may add them here.

*For more information on creating and using Custom Data Queries, see the **Custom Data Queries** section of this manual.*

Click 'Create' to save the Data Store and finish this wizard.



Queries: Data Queries ▾

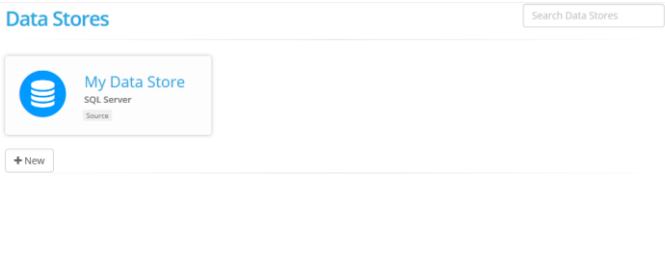
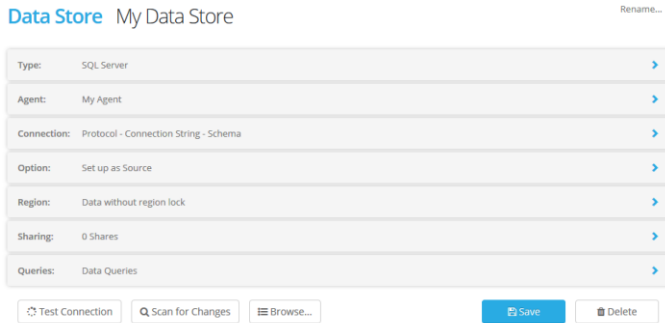
Would you like to build data queries? [Why would I need a Query?](#)

+ New

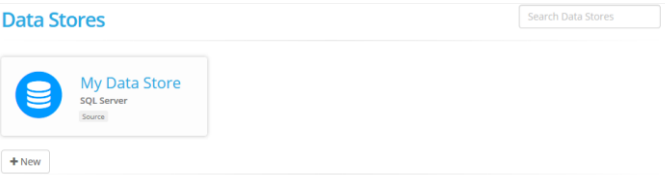
Create

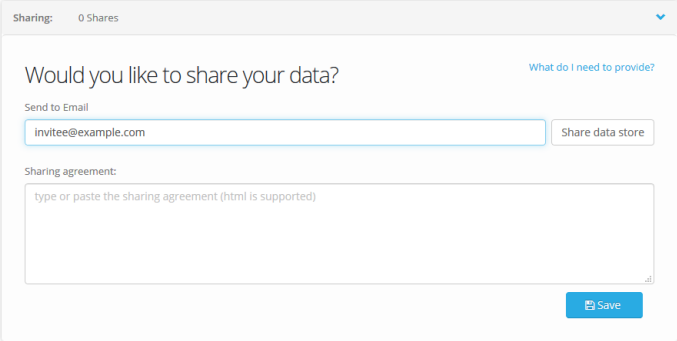
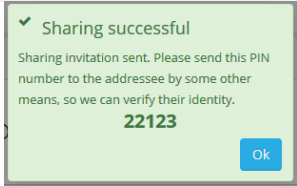
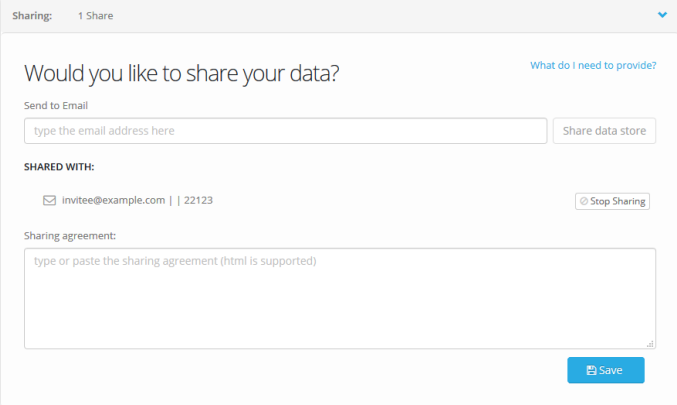
If you do not click the 'Create' button, your Data Store will not be created.

## Edit a Data Store

1.	To edit an existing Data Store, select it from the list of all Data Stores within the current Project.	 <p>The screenshot shows the 'Data Stores' page with a search bar at the top right. Below the header, there is a card for 'My Data Store' which is an 'SQL Server' source. A '+ New' button is visible below the card.</p>
2.	<p>The steps for editing a Data Store are the same as those for creating a Data Store.</p> <p><b>Please follow the steps described in the <i>Creating a Data Store</i> section of this manual.</b></p> <p>Click 'Save' to save all of your changes when you're finished.</p>	 <p>The screenshot shows the 'Data Store' edit page for 'My Data Store'. It has a 'Rename...' link at the top right. The main area contains a form with the following fields: Type (SQL Server), Agent (My Agent), Connection (Protocol - Connection String - Schema), Option (Set up as Source), Region (Data without region lock), Sharing (0 Shares), and Queries (Data Queries). At the bottom, there are buttons for 'Test Connection', 'Scan for Changes', 'Browse...', 'Save', and 'Delete'.</p>

## Share a Data Store

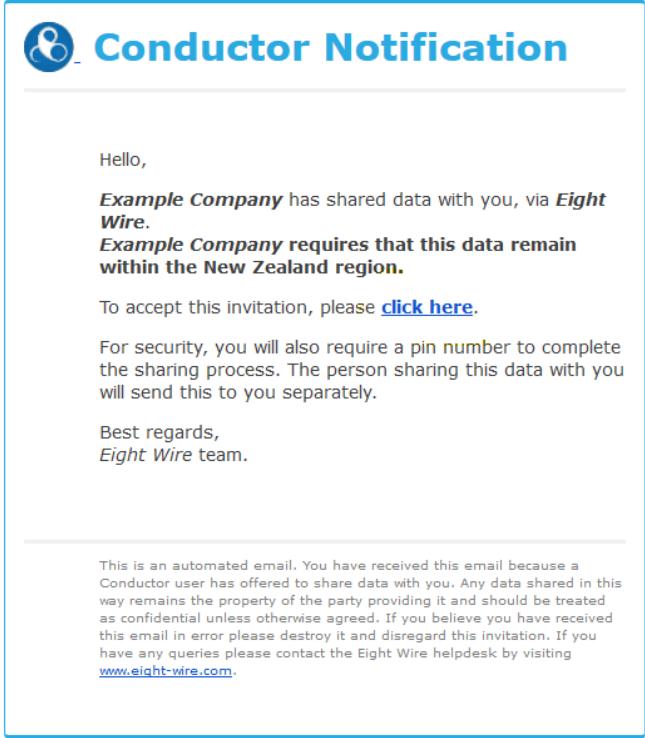
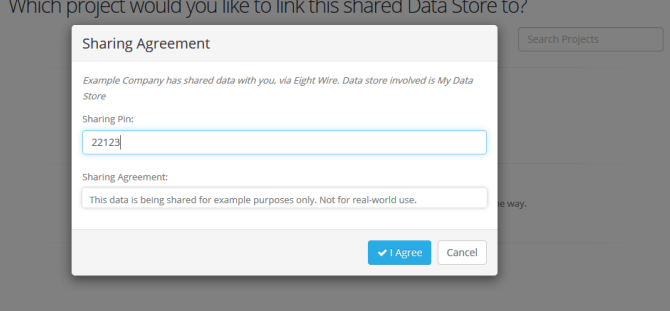
1.	<p>To share an existing Data Store, select it from the list of all Data Stores within the current Project.</p> <p>You can share a Data Store during its creation in much the same way – make sure you <b>save</b> the Data Store.</p>	 <p>The screenshot shows the 'Data Stores' page with a search bar at the top right. Below the header, there is a card for 'My Data Store' which is an 'SQL Server' source. A '+ New' button is visible below the card.</p>
----	---	--




<p>2.</p>	<p>Expand the <i>Sharing</i> panel for the selected Data Store.</p> <p>Enter the email address of the person you wish to share this Data Store with.</p> <p>Enter the reason this data is being shared. The recipient will be agreeing to this when they accept the invitation.</p> <p>Click 'Share data store' when you're finished.</p>	
<p>3.</p>	<p>The data sharing invitation email will be sent on your behalf and you will be shown a five-digit number. You must manually give this number to the data sharing recipient – it is not sent in the invitation email. This is for security reasons.</p> <p>Click 'OK' when finished.</p>	
<p>4.</p>	<p>If you forget the number, it will be shown on the Sharing panel for this Data Store until you stop sharing this Data Store.</p>	




## Accept a Data Sharing Invitation

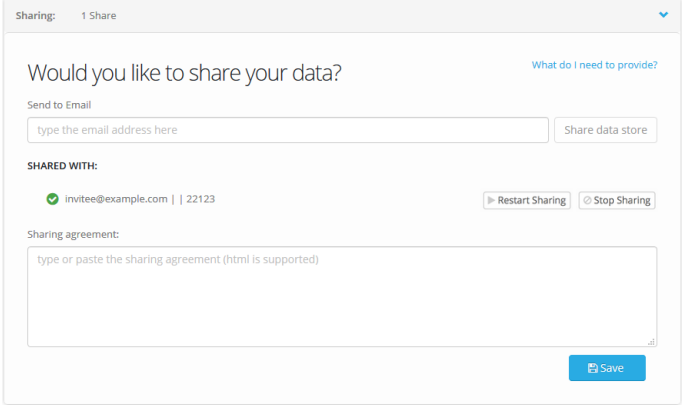
If you are sent an email entitled 'Eightwire – Invitation to Share Data', you have been invited to share external data. Follow these steps to start using this data.

<p>1.</p>	<p>If you have not already done so, log in to Eightwire.</p> <p>Now, open the data sharing email you received.</p> <p>This email will explain who has invited you to share data and if there are any data sovereignty restrictions.</p> <p>Click 'click here' in the email to go to the data sharing wizard in Eightwire.</p>	
<p>2.</p>	<p>When you click the link in the email, you will be taken to your web browser. If you have already logged into Eightwire, you will be asked for the Sharing PIN – this is a five-digit number that you must request from the sharing party – it will not be automatically sent to you, for security reasons.</p> <p>If you have the PIN, enter it now. If not, you can repeat this process later.</p> <p>If given, read the sharing agreement and if you agree, click 'I Agree' to accept the invitation.</p>	

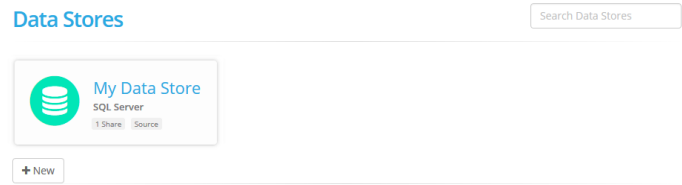
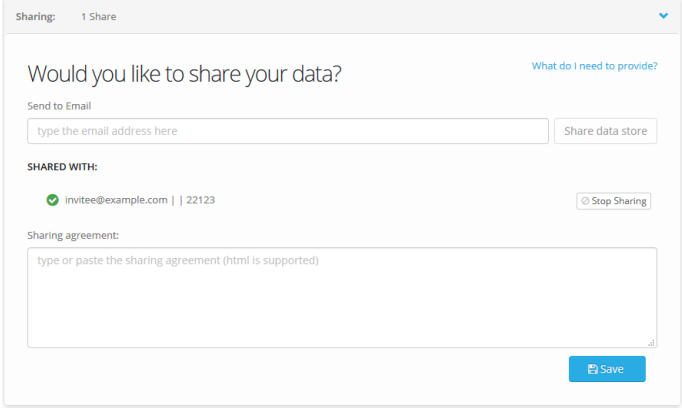
3.	<p>Choose which Project(s) this new shared Data Store will be created in by clicking the relevant 'Link' button(s).</p> <p><i>You can link this data share to more than one Project. If you do, one new shared Data Store will be created in each selected Project.</i></p>	<p>Which project would you like to link this shared Data Store to?</p> <p>Search Projects</p> <div> <div>  <p>Default Project...</p> <p>The default project.</p> <p>Link</p> </div> <div>  <p>My Project</p> <p>Link</p> </div> </div> <p>You may link this Data Store to as many projects as you wish. These Data Stores will look similar to any other Data Store and function in the same way. When you're finished here, you're ready to start using them.</p>
4.	<p>Now, you may use the shared Data Store like any other Data Store.</p> <p><i>You will not have access to the core properties of this Data Store because in sharing a Data Store the external party has granted you the use of the Data Store, but not the ability to view or modify its properties.</i></p>	<p>Data Stores</p> <p>Search Data Stores</p> <div>  <p>My Data Store</p> <p>SQL Server</p> <p>External Source</p> </div> <p>+ New</p>

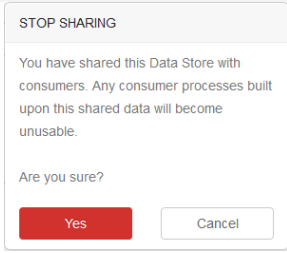
## Pause or Resume a Sharing Agreement

1.	<p>To pause an active sharing agreement or resume a paused sharing agreement, select a Data Store that you are sharing from the list of all Data Stores within the current Project.</p> <p><i>Shared data stores will have a label indicating that they are shared.</i></p>	<p>Data Stores</p> <p>Search Data Stores</p> <div>  <p>My Data Store</p> <p>SQL Server</p> <p>1 Share Source</p> </div> <p>+ New</p>
----	---	---

2.	<p>Expand the <i>Sharing</i> panel for the selected Data Store.</p> <p>Locate the share you wish to pause or resume and click either the 'Pause Sharing' or 'Resume Sharing' button next to it.</p>	
3.	<p>Confirm your intentions when prompted.</p>	

## Stop Sharing Data

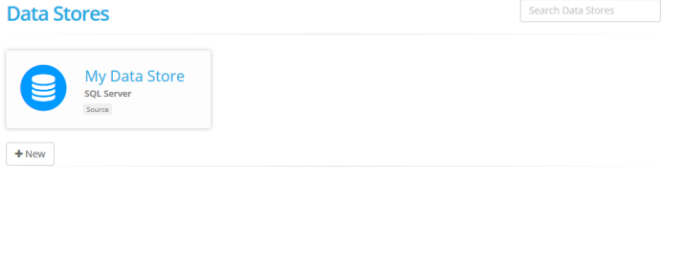
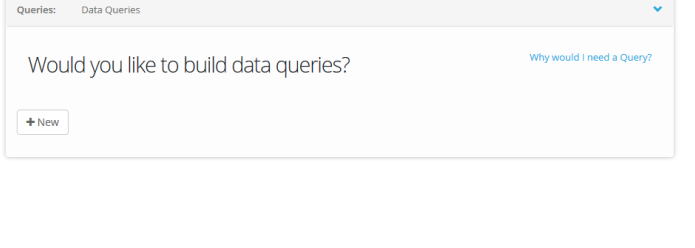
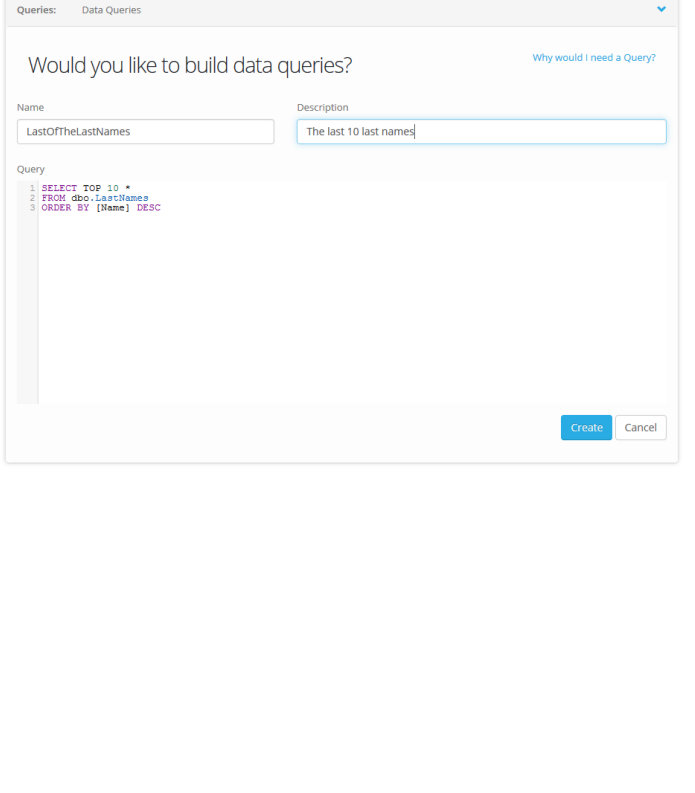
4.	<p>To stop sharing data, select a Data Store that you are sharing from the list of all Data Stores within the current Project.</p> <p><i>Shared data stores will have a label indicating that they are shared.</i></p>	
5.	<p>Expand the <i>Sharing</i> panel for the selected Data Store.</p> <p>Locate the share you wish to cancel and click 'Stop Sharing' next to it.</p>	

6.	<p>Confirm that you wish to stop sharing this Data Store by clicking 'Yes'.</p> <p><i>Stopping sharing will permanently delete this data sharing agreement. If you wish to temporarily stop sharing, pause the share instead.</i></p>	
----	---	--

**Any Processes that depend on this Data Store, including those belonging to the recipient, will break or be deleted when sharing is stopped. You should confirm with the recipient before doing this.**

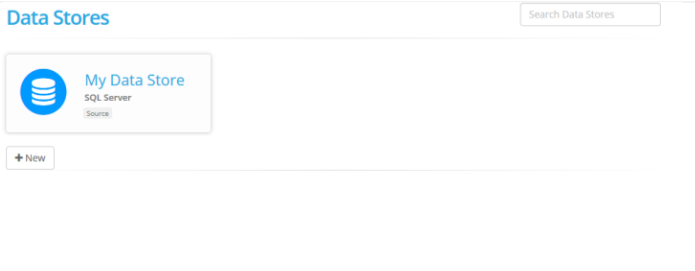
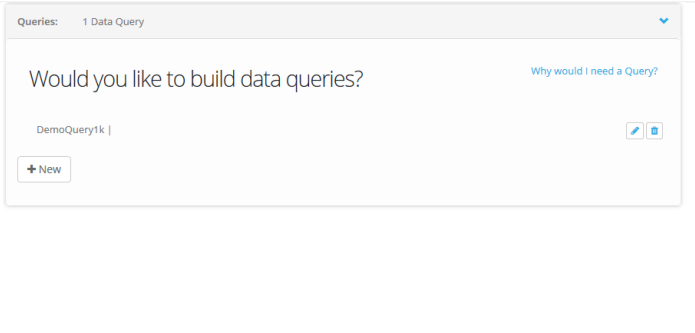
## Create a Custom Data Query

Only Data Stores connecting to relational databases support Custom Data Queries.

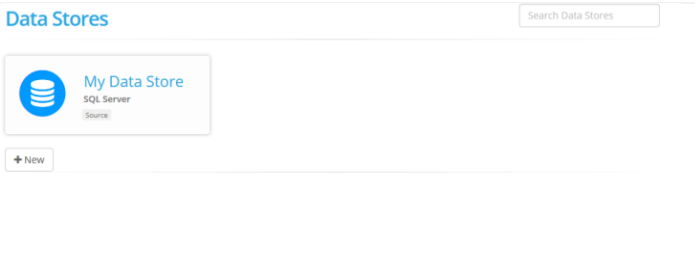
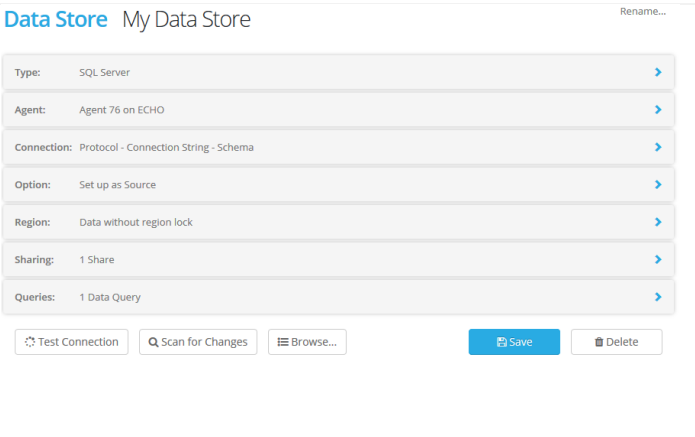
1.	To create a Custom Data Query, select a Data Store from the list of all Data Stores within the current Project.	
2.	Expand the Data Store's 'Queries' panel.  Click 'New' to create a new query.	
3.	<p>Enter a name and a description for your query.</p> <p>Enter the SQL for your query. This SQL will be executed against the Data Store's underlying database and the data returned will be used by Eightwire as if it came from a table.</p> <p><i>You will be able to use this query when Creating Processes as if it was a table.</i></p> <p>Click 'Create' to save the query.</p>	

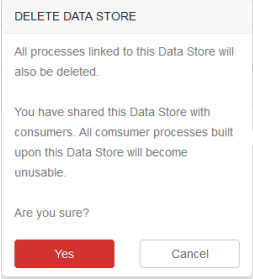
## Edit or Delete a Custom Data Query

*Only Data Stores connecting to relational databases support Custom Data Queries.*

1.	To delete a Custom Data Query, select a Data Store from the list of all Data Stores within the current Project.	
2.	<p>Expand the Data Store's 'Queries' panel.</p> <p>To edit or delete existing queries, click the 'Edit' or 'Delete' icons next to the query.</p>	

## Delete a Data Store

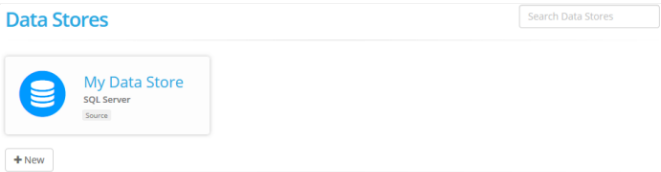
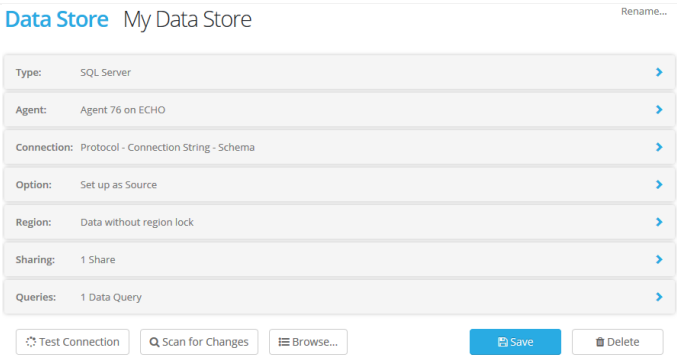
1.	To delete a Data Store, select the Data Store from the list of all Data Stores within the current Project.	
2.	Click 'Delete' at the bottom of the page.	

3.	When asked to confirm the deletion, click 'Yes'.	 <p><b>DELETE DATA STORE</b></p> <p>All processes linked to this Data Store will also be deleted.</p> <p>You have shared this Data Store with consumers. All consumer processes built upon this Data Store will become unusable.</p> <p>Are you sure?</p> <p><b>Yes</b> Cancel</p>
----	--	--

**Deleting a Data Store will also delete all Processes, Custom Data Queries and Filters you may have created with it. If the Data Store is shared, this will also impact the sharing party.**

### Browse and Edit Data Structures

When you create a Data Store, Eightwire will scan the existing data structures in that Data Store and remember them. For a database, this means Eightwire will gather structural information about those database tables. For a CSV file, this means Eightwire will open each CSV file in the target folder and sample the data, inferring sensible data types for each column. This remembered structure information is available for you to browse and modify.

1.	To browse a Data Store, select the Data Store from the list of all Data Stores within the current Project.	 <p><b>Data Stores</b> Search Data Stores</p> <p>My Data Store SQL Server Source</p> <p>+ New</p>
2.	At the bottom of the Data Store details page, click 'Browse'.	 <p><b>Data Store</b> My Data Store Rename...</p> <p>Type: SQL Server &gt;</p> <p>Agent: Agent 76 on ECHO &gt;</p> <p>Connection: Protocol - Connection String - Schema &gt;</p> <p>Option: Set up as Source &gt;</p> <p>Region: Data without region lock &gt;</p> <p>Sharing: 1 Share &gt;</p> <p>Queries: 1 Data Query &gt;</p> <p>Test Connection Scan for Changes Browse... Save Delete</p>

3. This page lists all of the tables that Eightwire is able to see in the Data Store. Click on one to expand it.

**Browsing** My Data Store  
This how Conductor sees your data store

- > LastOfTheLastNames (data query)
- > dbo.FemaleFirstNames
- > dbo.LastNames
- > dbo.MaleFirstNames
- > dbo.ObfuscationTestData
- > dbo.Standard100K
- > dbo.Standard10K
- > dbo.Standard1K
- > dbo.Standard1K\_A
- > dbo.Standard1K\_B
- > dbo.Standard1K\_C
- > dbo.Standard1K\_D
- > dbo.Standard1K\_E
- > dbo.Standard1K\_F

4. Here, you will see the table structure – one row for each column. This panel shows each columns name, data type, if it is part of a primary key, whether it allows 'null' values and whether it is read-only (not writable).

To ignore the column and not use it, uncheck the 'Use' checkbox on the left.

To edit any of these properties, click the edit icon to the right.

To delete a column, click the delete icon to the right.

To undo any *unsaved* changes, click the undo icon to the right.

**When you are finished making changes, click 'Save'.**

To add a new column, click 'New Column'.

To rescan the underlying data store for changes, click 'Rescan' – **any changes made manually will be lost.**

> LastOfTheLastNames (data query)  
> dbo.FemaleFirstNames  
▼ **dbo.LastNames**

Structure Filter Tags

Use	Name	Type	Key	Nullable	Read-Only	Is Unique	Unicode	
<input checked="" type="checkbox"/>	ID	NUMBER (10,0)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="button" value="edit"/> <input type="button" value="delete"/> <input type="button" value="undo"/>
<input checked="" type="checkbox"/>	Name	TEXT (50)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="button" value="edit"/> <input type="button" value="delete"/> <input type="button" value="undo"/>

Use

☒ ☐ ID

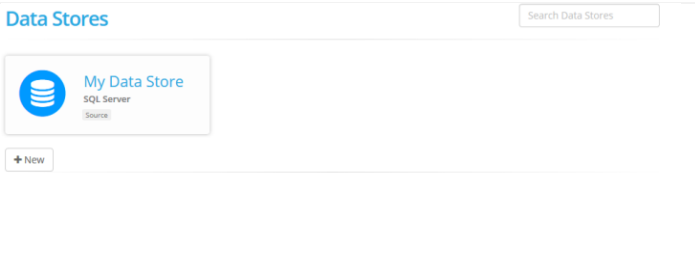
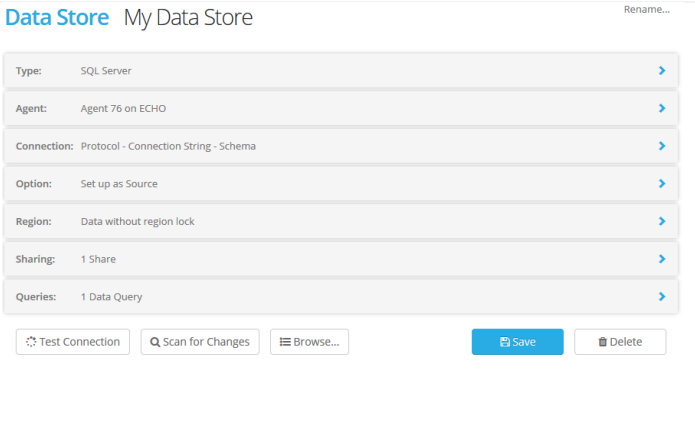
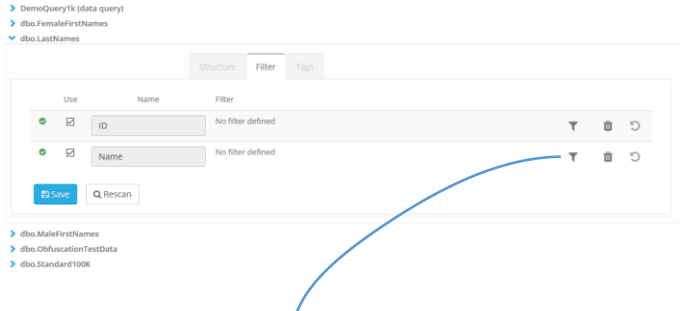

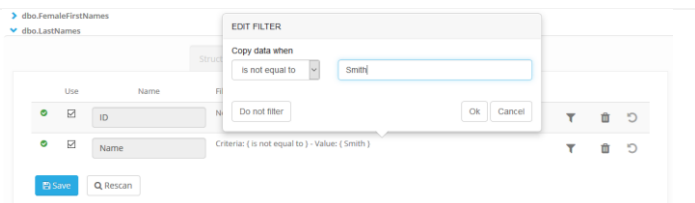
Edit column

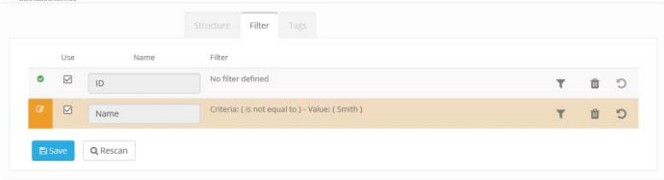
Delete Column

Undo changes

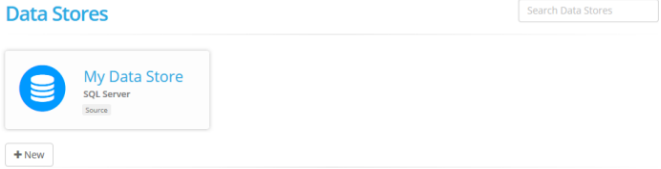
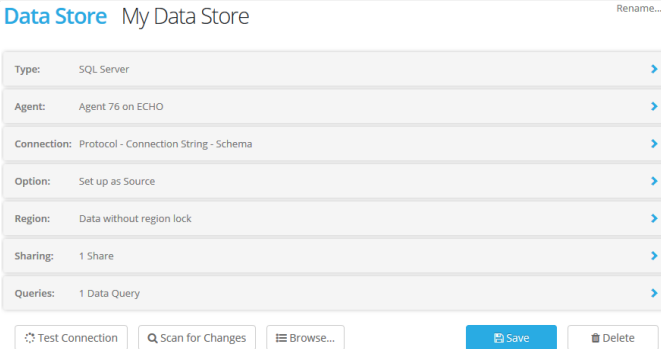
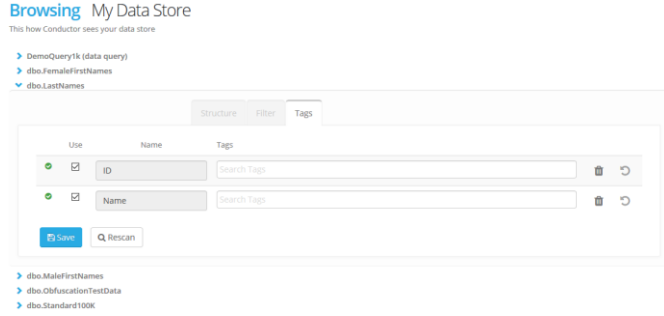


## Manage Data Store Filters

1.	To manage the row-level filtering for tables in a Data Store, select the Data Store from the list of all Data Stores within the current Project.	
2.	At the bottom of the Data Store details page, click 'Browse'.	
3.	Select the table of interest, and click on it to expand it.  Click the 'Filter' tab to show the column filters (if any) for this table.	
4.	Click the filter icon next to the column you wish to edit.	
5.	This popup dialog allows you to edit the filter criteria for this column.  <i>Only rows where all columns meet the filter criteria (if filtered) will be used by Eightwire. This means that for a source Data Store, only matching rows will be read; for a destination Data Store, only matching rows will be written.</i>	 <p>In this example, only rows where the <b>Name is not 'Smith'</b> will be included.</p>

	To remove this column's filter, click 'Do not filter'.	
6.	<p>To <b>undo</b> any <b>unsaved</b> changes (shaded), click the undo icon to the right of each column.</p> <p><b>When you are finished making changes, click 'Save'.</b></p>	


## Manage Data Store Tags

1.	To manage the tags applied to Data Store structures, select the Data Store from the list of all Data Stores within the current Project.	
2.	At the bottom of the Data Store details page, click 'Browse'.	
3.	<p>Select the table of interest, click on it to expand it.</p> <p>Click the 'Tags' tab to show the currently applied column tags (if any) for this table.</p>	

4. To add a tag, simply click into the blank part of the Tags field for a column and start typing. Any matching tags will be showing – select one to add it.


To add more tags, type the name of the next tag, select it, and so on.

To remove a tag, click the delete icon next to the tag's name.



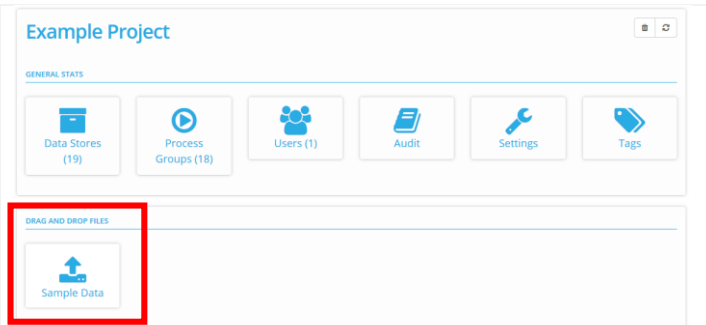
5. When you are finished, click **'Save'**.

*See the 'Tagging' section of this User Guide to learn more about how to manage the list of available tags.*




## Transfer Data via *Drag and Drop*

1. If a Drag and Drop Data Store has been shared with you or you have created one and connected it to a destination using Processes, locate the Drag and Drop area of your Project Dashboard.




2. Locate a file that is appropriate for the type of Drag and Drop Data Store you are using (either Delimited Text (CSV) or Excel Spreadsheet). If you aren't sure how to do this, click on the Drag and Drop area to show a dialog box that will help you locate a file.

*Confirm that you want to upload this file:*



*You will see this message while your file is being uploaded:*



Upload that file, either by dragging it from its location onto the Drag and Drop area of the web portal, or by selecting it in the dialog box. This will automatically begin the data transfer process.

If there is a problem, you will be notified, but you may need to contact the receiving organisation for more details if this Data Store was shared with you.

## Agents

Eightwire is a cloud service - it operates in the internet. Many companies store their data safely inside their own corporate network, making it inaccessible from the internet. In order for Eightwire to transfer data to or from these corporate data stores, the data needs to be accessible to Eightwire. The most secure way to achieve this access is via Agents.

An Agent is a small piece of software that operates inside your corporate firewall, usually on your corporate network. The Agent communicates securely with the Eightwire cloud service and the Agent is able to read and write the data within your corporate network that it has been given permission to access. The Agent becomes Eightwire's 'eyes and ears' inside your corporate network.

An Agent does not usually require any modification to your firewall rules in order to operate and is very easy to install and get connected.

### Secure Communications

Each Agent establishes a secure TLS 1.2 connection to Eightwire's cloud services. All data sent over this connection is securely encrypted as it travels to and from Eightwire's cloud services. Agents only make outbound internet connections to Eightwire and do not allow incoming connections – this is a very secure design with a very low 'attack surface' and as such makes it one of the most secure methods of internet data transfer available.

### Direct vs Agent Connections

When configuring Data Stores in Eightwire, you will need to choose whether or not to use an Agent. If your data is stored 'in the cloud', you may not need an Agent – Eightwire's cloud services should be able to access it directly. In this scenario, the connection details you give to Eightwire should make sense from the point of view of Eightwire's cloud services.

If your data is inside a corporate network or behind a firewall, you should consider using an Agent. In this scenario, the connection to your data is made from the Agent on behalf of Eightwire, so your connection details will be from the Agent's perspective, not from the perspective of the Eightwire cloud services.

## System Requirements

The following requirements should be considered before installing an Agent.

### Host Computer

#### Minimum Specification:

CPU	2 Core
RAM	6 Gb
Disk	100 Gb HDD
OS	Windows 7 (with .NET Framework 4.5 installed)
	Windows 8
	Windows Server 2012

#### Recommended Specification:

CPU	4 Core
RAM	8 Gb
Disk	100 Gb SSD
OS	Windows 10
	Windows Server 2016

Agents do not store data on disk and only write the occasional error log and monitoring/update logs, so a small disk is usually adequate. As Agents become busier, they will consume more CPU and RAM. Increasing CPU and RAM will reduce the chance of the Agent becoming a bottleneck, however most organisations will encounter internet bandwidth constraints well before they encounter an Agent bottleneck.

### Permissions

The Agent runs as a Windows Service. By default, the service runs under the Local System (NT AUTHORITY\System) user account. You can configure this either during installation or by changing the service settings and restarting it.

Unless you only intend for the Agent to access resources on the same computer as the Agent, it is a good idea to create a new service account in your networks Active Directory, specifically for the purpose of running the Agent. Set the Agent service to run under this new account.

You should then restrict the access this service account has on your network to only allow access to those resources you wish to use with Eightwire. For example, give read access to a folder containing CSV files you want to use as a source, or give write permissions on a SQL Server database you wish to write data into using Eightwire (and use a trusted connection string in Eightwire).

Securing database access (on those platforms that support integrated Windows security) in this way means that you do not need to include a username and password in the connection string given to Eightwire, which is more secure and good practice.

Creating a new domain account for each Agent means that each Agent's access can be specifically controlled by you or revoked entirely if required.

In addition to file and database access, this service account also needs modify (read/write/create) access to the Windows RSA Key Store (C:\ProgramData\Microsoft\Crypto\RSA\MachineKeys). This is required to allow the Agent to encrypt the contents of its own configuration file to ensure a high level of security.

## Internet Access

The Agent requires limited internet access to communicate with Eightwire's cloud services. The Agent will attempt to establish secure HTTPS connections on port 443. These connections are all outbound from the Agent and will always be to services on the *eight-wire.com* domain.

Establishing internet access is often the most difficult part of installing an Agent within an enterprise network. Network administrators need to allow internet access from this service account, which is sometimes overlooked.

The local or domain service account under which the Agent is running requires either direct internet access, or to be able to automatically detect and connect to your corporate proxy server, or have manual proxy settings created during installation or subsequently through the Agent Configuration Utility.

If the service account is created with the correct profile information, the Agent should be able to automatically detect and use your corporate proxy server if you set the Agent to auto-detect a proxy server. If this doesn't work, you can manually enter proxy credentials during Agent installation or later through the Configuration Utility. If this is unsuccessful, you may need to add proxy bypass rules for this service account or computer, allowing the Agent access to internet services on the *eight-wire.com* domain (e.g.: \*.*eight-wire.com*).

Another consideration for network administrators is the need for DNS resolution. The Agent connects to Eightwire services based on fully qualified domain names, not IP addresses. This means that DNS resolution is required. If the Agent is installed in a DMZ or other secure area that does not allow DNS resolution, then this poses a challenge to running the Agent there. The Eightwire service maintains data sovereignty and manages its own load-balancing by being able to add or remove services which are internet facing from the cloud. Therefore it is necessary that each Agent be able to resolve and connect to any number of Eightwire services as required. *Hard-wiring IP addresses in a host's file to avoid this requirement is not good practice.*

## Data Access

It may seem obvious, but the Agent requires access to your data, otherwise it can't do its job. This can be more difficult than it sounds. As mentioned previously, you should allow specific permissions to the data that you want the Agent to work with and deny access to anything you don't want the Agent to be able to access.

Remember, you control what the Agent accesses through the web portal anyway, but as a second-level of protection it is good practice to only allow access to resources that you want access to and broadly deny access to everything else.

One common approach to installing the Agent within enterprise networks is to install it on a computer within a DMZ. A DMZ is a 'no-mans-land' sandbox network between the internet and your company's core internal networks. This is seen as a 'safe' place to install the Agent because it is specifically blocked from the internal network. However, the Agent needs access to resources on the internal network. Often there is no domain trust between the internal and DNZ networks, further complicating things. If you must use this approach, you will need to do one of the following to make it work:

1. Allow the Agent to communicate inwards through the internal firewall to specific resources and either trust it or have it pass credentials through (less secure). This allows for less manual data shipping, but requires some care to ensure the internal firewall rules allow only the access required and nothing more.
2. Create a 'data staging' area inside the DMZ and have data shipped between this staging area and your internal network by some other trusted process. You would then give the Agent access to this data. This is notionally more secure, but requires significant manual work to move data around, which is usually what companies were trying to avoid by using Eightwire in the first place.

Installing the Agent on a well-secured computer within your internal network is arguably simpler and for most small organisations, will be acceptable.

These are some of the considerations that should be taken into account when deciding where to install an Agent.

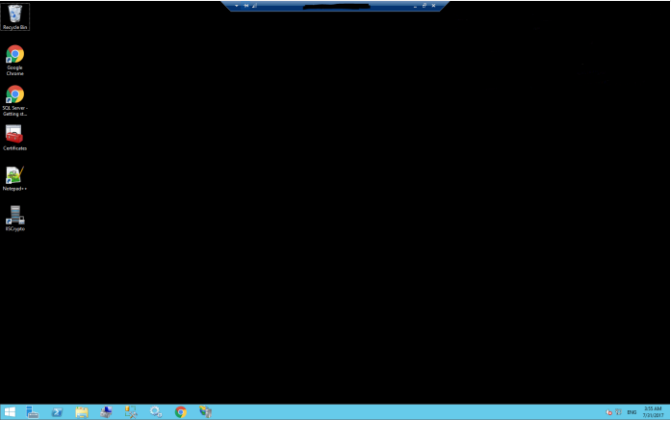



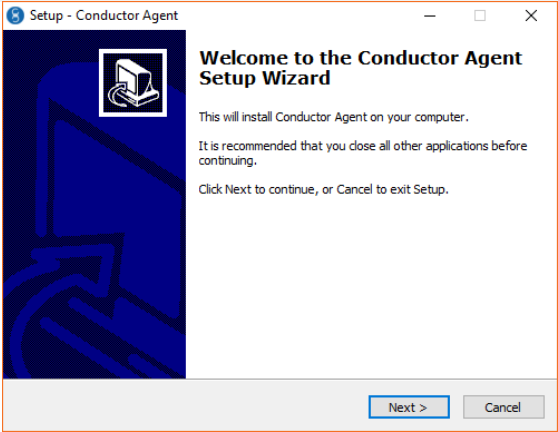
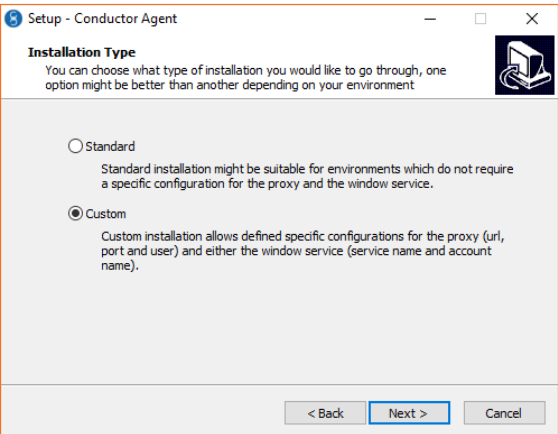
## Tasks

These should all be performed directly on the server or PC where the Agent is installed.

### Install an Agent

There are two situations where you might be installing an Agent – either you have just created a new Data Store and have downloaded a new Agent Installer as part of this wizard, or you are reinstalling an existing Agent.

<p>1.</p>	<p>Log onto the Windows computer where you wish to install the Agent.</p> <p>If you have the Agent installer file, continue to the next step. If you do not, follow either the <i>Create a Data Store</i> task or the <i>Download an Existing Agent Installer</i> task in the manual in order to download an Agent installer.</p>	
<p><b>You must have <i>local administrator</i> privileges on this computer to install an Agent.</b></p>		
<p>2.</p>	<p>Locate the downloaded Agent installation file and execute it.</p> <p>When you execute the Agent installer, you may be asked to confirm that you want to allow this app to make changes. Click 'Yes' to continue.</p> <p><i>The installer requires administrative access to create the Windows Service.</i></p>	

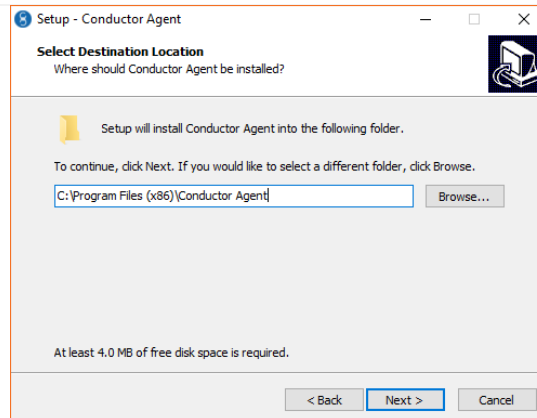
3.	<p>The installer will open.</p> <p>Click 'Next'.</p>	
4.	<p>Choose either the 'Standard' or 'Custom' installation. If you choose 'Standard', the most common settings will be chosen for you to make the installation quicker.</p> <p><i>In this example, we will choose 'Custom' to show you all of the options available.</i></p> <p>Click 'Next'.</p>	

5. Choose where to install the Agent. If you wish to install multiple Agents on this computer, you will need to install them into different paths.

The path you select **must not** contain an existing running Agent. To replace an existing Agent, properly stop the service first.

*The name of the child folder in the path will become the name of the Windows Service.*

Click 'Next'.

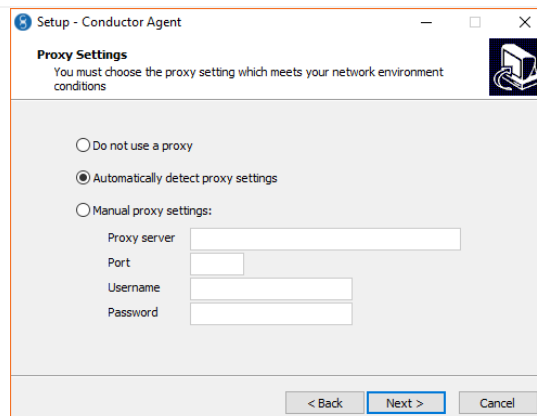


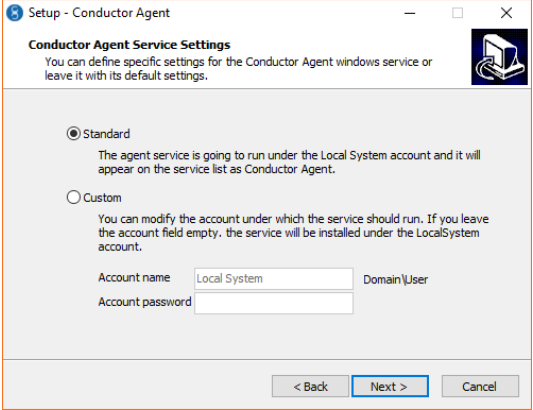
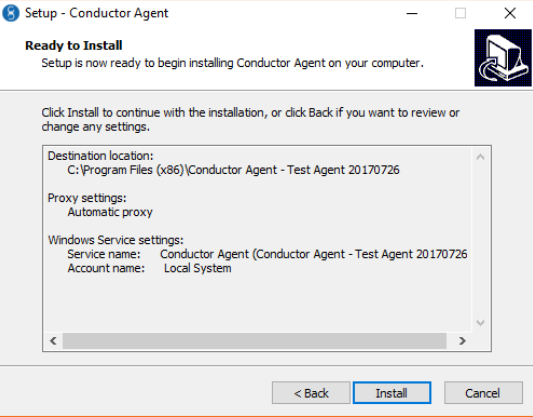
6. If your organisation uses a proxy server to access the internet, select either 'Automatic' or 'Manual'.

If automatic detection doesn't work, try manually setting your proxy settings. *Your proxy settings will be encrypted and only be accessible to the Agent.*

If you do not need to use a proxy server, select 'Do not use a proxy'.

Click 'Next'.



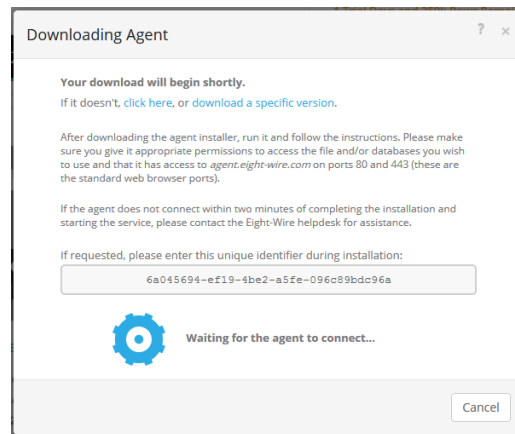
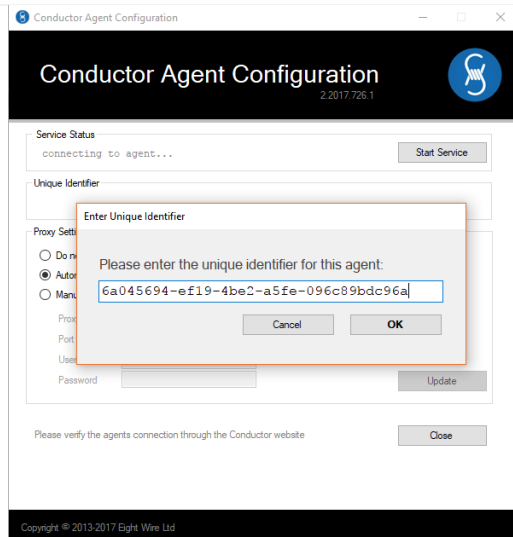
7.	<p>Choose which Windows security account the Windows Service will run under.</p> <p><i>By default the Windows Service runs as</i></p> <p><i>NT AUTHORITY\System.</i></p> <p>Click 'Next'.</p>	
8.	<p>The installer has all the information it needs to begin the installation. Check the information is correct and when you're ready, click 'Next'.</p>	

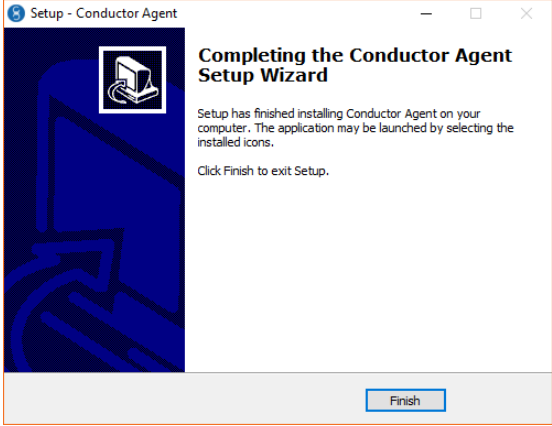
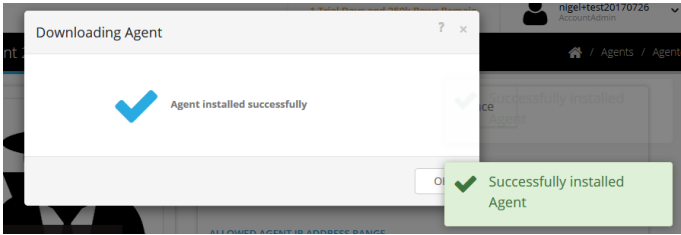

9. The installer will try to detect the Agent's unique identifier from the Agent's installer file, but if it can't, you will be asked to enter the unique identifier.

The unique identifier will be shown on the web page you downloaded the installer from, if that page is waiting for you to complete the installation.

Here, you can see the web page, waiting for the Agent to connect. This page shows the unique identifier and instructs you to provide it to the installer if requested.


Click 'Ok' to complete the installation (if you are asked to complete this step).



10.	<p>When the installation is complete, you should be able to verify that the Agent has successfully connected to Eightwire by checking the web page that was waiting for the installation to complete.</p>	 
11.	<p>If you are not sure how to check if the Agent is connected or you are not looking at one of the web pages mentioned, navigate to the Agents section of the website and look for the Agent you just installed. You should see a green tick mark next to the Agent.</p>	

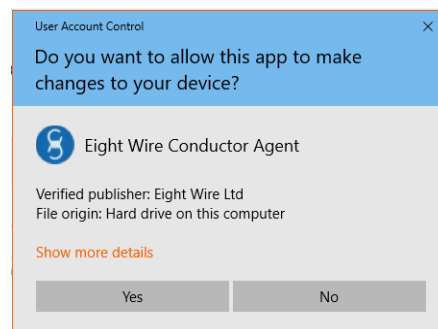
## Configuring an Agent

1. Log onto the Windows computer where the Agent is installed.  
  
Locate the folder where the Agent is installed. This will likely be:  
  
*C:\Program Files (x86)\Conductor Agent*  
  
Run the 'AgentConfiguration.exe' utility.



**You must have *local administrator* privileges on this computer to configure an Agent.**

2. You may be asked to confirm that you want to allow this app to make changes. Click 'Yes' to continue.  
  
*The Agent Configuration Utility requires administrative access to control the Windows Service.*

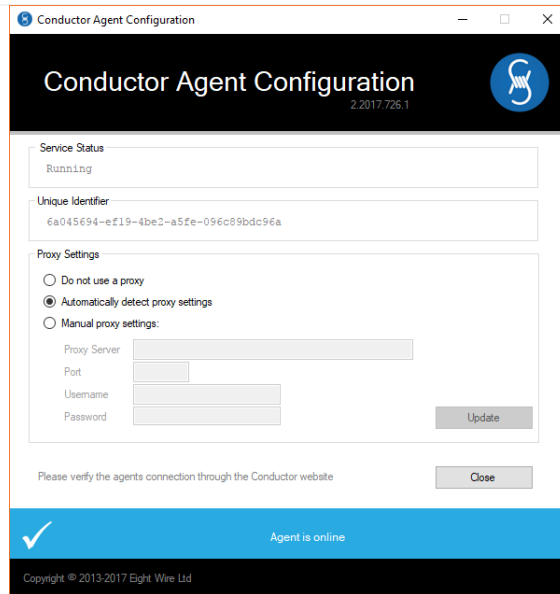


3. This utility will tell you if the Agent service is running or not, what its unique identifier is and what proxy settings in use.

If the Agent is not online, the error will be shown here and may be useful in diagnosing the problem.

When you are finished, click 'Close'.

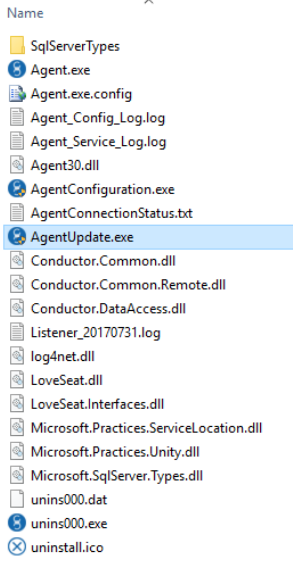

*If you make changes to the proxy settings, it may be necessary to restart the Agent Windows service for the changes to take effect.*





## Manually Updating an Agent

If you have configured your Agent to **not** automatically update itself, you will have to perform manual updates from time to time – usually when advised to do so by Eightwire support staff.

<p>1.</p>	<p>Log onto the Windows computer where the Agent is installed.</p> <p>Locate the folder where the Agent is installed. This will likely be:</p> <p><i>C:\Program Files (x86)\Conductor Agent</i></p> <p>Run the 'AgentUpdate.exe' utility.</p>	
<p><b>You must have <i>local administrator</i> privileges on this computer to update an Agent.</b></p>		
<p>2.</p>	<p>You may be asked to confirm that you want to allow this app to make changes. Click 'Yes' to continue.</p> <p><i>The Agent Update Utility requires administrative access to control the Windows Service and replace files.</i></p>	

3. The update utility runs on the command line and will show a command console output similar to the one shown here.

There is a lot of information, but in particular you should scroll to the top and look for any errors.

Usually you will either see a list of files being updated or you will see a message saying that the Agent is up-to-date. The information displayed after this is related to the re-encryption of the Agent's configuration file and can be ignored.

If there were any problems with the update, all changes will be rolled back and the unmodified Agent service restarted.

Press any key to exit close the update utility.

```

C:\Program Files (x86)\Conductor Agent - Test Agent 20170726\AgentUpdate.exe
Conductor Agent Update
Copyright © 2013-2017 Eight Wire Ltd

Update Requested
Update Requested for Agent installed on: C:\Program Files (x86)\Conductor Agent
- Test Agent 20170726\Agent.exe
Agent installed version: 2.2017.726.1
Agent installed path: C:\Program Files (x86)\Conductor Agent - Test Agent 201707
26
Agent installed guid: 6a045694-ef19-4be2-a5fe-096c89bdc96a
Set proxy: automatic
Current version is 2.2017.726.1
Checking for new version...
No updates are available
The Agent is up-to-date.
Microsoft Windows [Version 10.0.15063]
(c) 2017 Microsoft Corporation. All rights reserved.

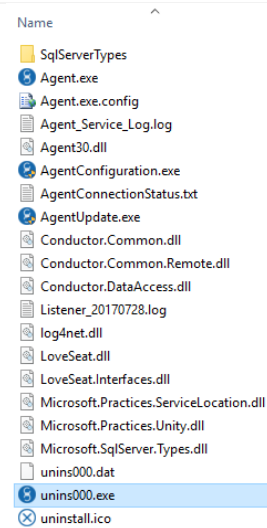
C:\Program Files (x86)\Conductor Agent - Test Agent 20170726>cd %windir%\Microso
ft.NET\Framework\v4.0.30319
C:\Windows\Microsoft.NET\Framework\v4.0.30319>rename "C:\Program Files (x86)\Con
ductor Agent - Test Agent 20170726\Agent.exe.config" web.config

```

**Forcing a manual update could interrupt data transfers already in progress.  
The Agent Service will be stopped during the update and restarted afterwards.**

## Uninstalling an Agent

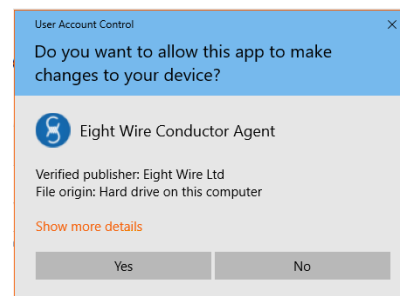
1. Log onto the Windows computer where the Agent is installed.
- Locate the folder where the Agent is installed. This will likely be:
- C:\Program Files (x86)\Conductor Agent*
- Run the 'unins000.exe' uninstall utility.
- Alternatively you can uninstall from the Windows Start Menu or from Control Panel > Programs and Features.*




**You must have *local administrator* privileges on this computer to uninstall an Agent.**

2. You may be asked to confirm that you want to allow this app to make changes. Click 'Yes' to continue.

*The uninstaller requires administrative access to stop the Windows Service and remove files.*



3. When prompted, click 'Yes' to begin the uninstallation process.

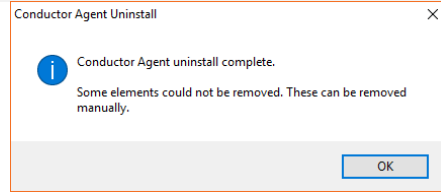


4.

When the uninstaller completes, click 'OK'.

*At this point the Windows Service has been removed and the Agent has stopped.*

*If the Agent was particularly busy when it was shut down, it is possible that some files were not removed properly. It is a good idea to check the installation folder and remove any remaining files.*



## Process Groups

### About Process Groups

A single Process is concerned with transferring the data in a single table. A Process Group is a collection of Processes that are executed together. For example, if you needed to transfer all of the Human Resources data from a larger database, you might group all of these related Processes into one Process Group.

When a Process Group is executed, Eightwire will attempt to orchestrate the data transfer of all tables in this group in such a way that does not break any of the referential constraints between these tables. This mainly applies to relational databases, but is important when dealing with tables that have relationships between them. Individual Processes can be executed if desired.

A Process Group can be configured to execute on a schedule. Schedules can be either one-off or recurring.

Each Process must belong to a Process Group.

### Destination Actions

Each Process Group has a configured Destination Action. This behaviour applies to all Processes in a Process Group. The Destination Action describes how each data transfer will write data to its destination. Not all actions are available for all types of Data Store, so some care is required to ensure you do not create Processes which write data into a Data Store that does not support the destination action you have selected.

#### Overwrite existing data

Any residual data in the destination table is deleted, then all source data is written to it. *Available for all types of Data Store.*

#### Append to existing data

All source data is added to the end of any existing data in the destination table.

#### Merge new and changed data

Primary keys are compared between existing source and destination data. Based on primary key data, any rows that do not exist in the destination are appended; any existing data is updated. *Available for relational database Data Stores only.*

#### Merge all data (insert/update/delete)

Primary keys are compared between existing source and destination data. Based on primary key data, any rows that do not exist in the destination are appended; any existing data is updated; any rows that exist in the destination but not in the source are deleted from the destination. *Available for relational database Data Stores only.*

#### Incrementally load data


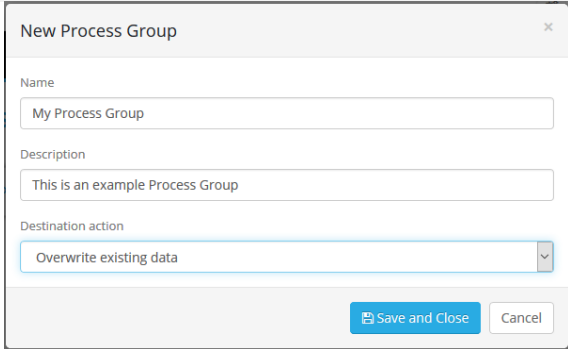
For each Process, an incremental column is nominated (see *Editing a Process*). This column must be a scalar value, such as a number or a date, and the data in the source table should only ever be added to, not updated. When data is transferred, the largest value in the incremental column in the destination table is taken as the starting point for reading source data – only rows with a larger value in the source will be transferred and appended to the destination table. Good candidates for this type of action are log or transaction tables. This is the most efficient method of transferring recently

added data from large tables provided the data is not updated after it has been added to the source table. *Available for relational database Data Stores only.*

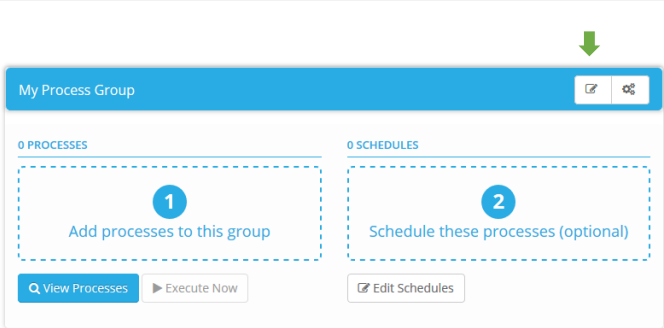
## Tasks

To create, edit or delete Process Groups, you need to be on the **Process Groups List** page of a Project. From the Dashboard, navigate to *Projects*, choose one Project from the list, and then navigate to *Process Groups* (labelled *Processes* on the navigation toolbar or *Process Groups* on the Project Dashboard icon).

### Create a New Process Group

1.	Click 'New' to show the <i>New Process Group</i> dialog.	
2.	<p>Enter the name and (optionally) the description of your Process Group.</p> <p>Choose the destination action for all Processes in this group.</p> <p>Click 'Save' when finished.</p>	

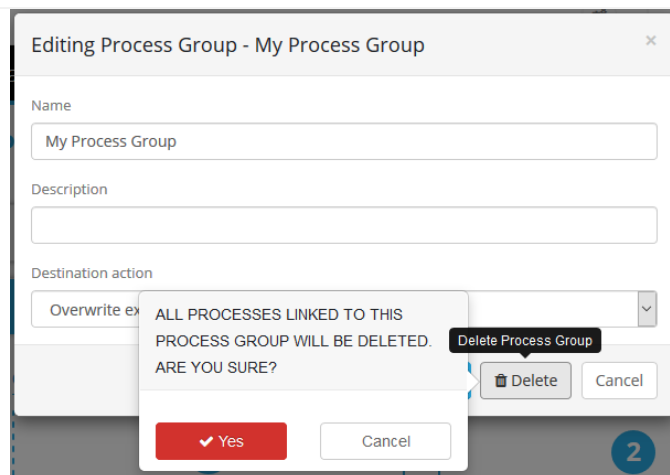
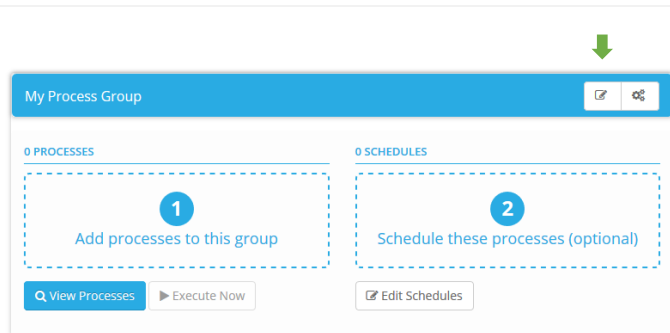
### Edit a Process Group

1.	Locate the Process Group in the list and click the small edit button on the right of the group title bar.	
----	---	--

2. Edit the name and (optionally) the description of your Process Group.  
  
Choose the destination action for all Processes in this group.  
  
Click 'Save' when finished.

### Delete a Process Group

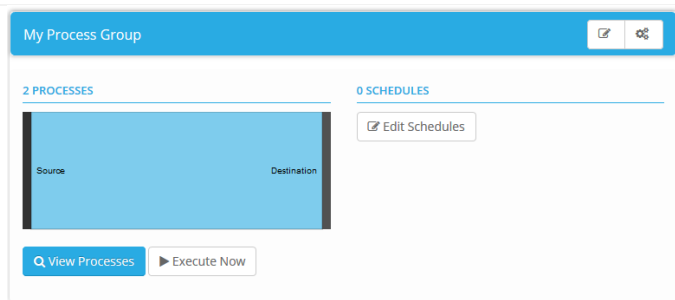
1. Locate the Process Group in the list and click the small edit button on the right of the group title bar.
2. Click 'Delete'.  
  
Click 'Yes' to confirm the deletion.



## Execute a Process Group

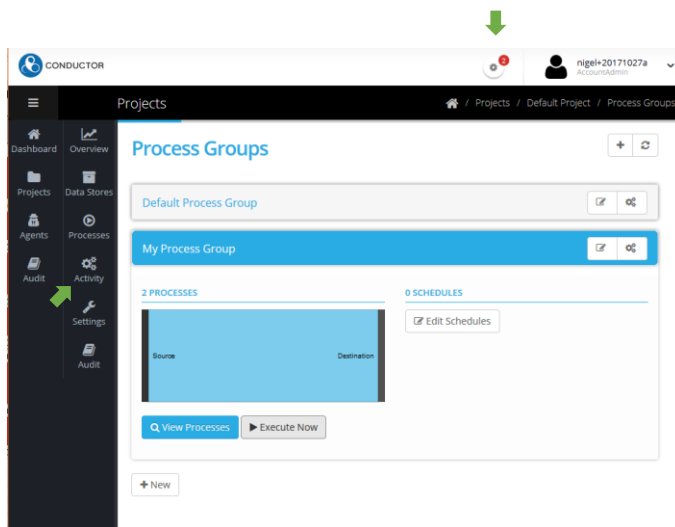
1. Locate the Process Group in the list and expand it.

Click 'Execute Now' to immediately execute all Processes in this Group.



2. After a few moments, the Processes will start executing in the background.

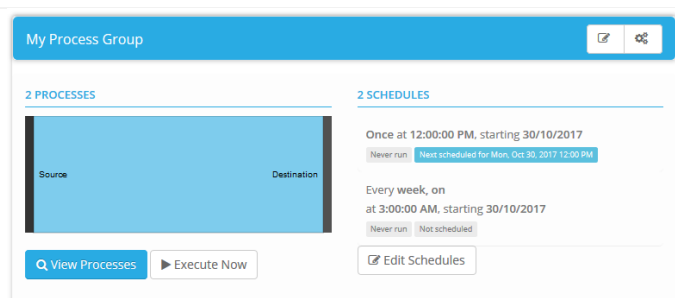
You should notice the activity indicator at the top of the screen animating and showing some Process activity. You can click on this to expand this mini-activity panel, or you can click 'Activity' in the left-hand navigation bar to view the main activity page if you wish to monitor the progress of these Processes.



## Schedule a Process Group

1. Locate the Process Group in the list and expand it.

Click 'Edit Schedules'.

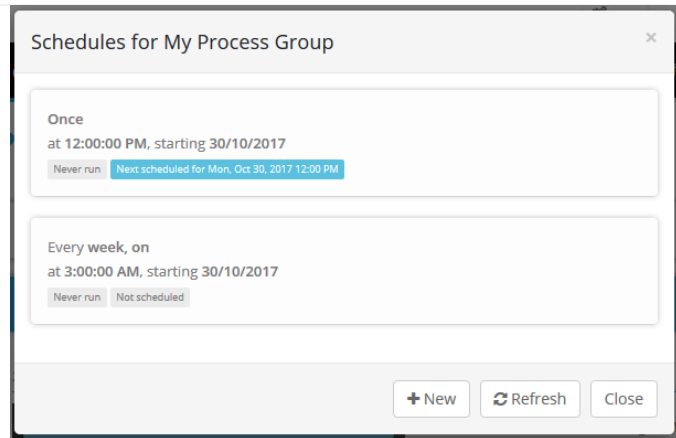




2. To create a new schedule, click 'New'.

*or*

To edit or delete an existing schedule, click on the schedule itself.

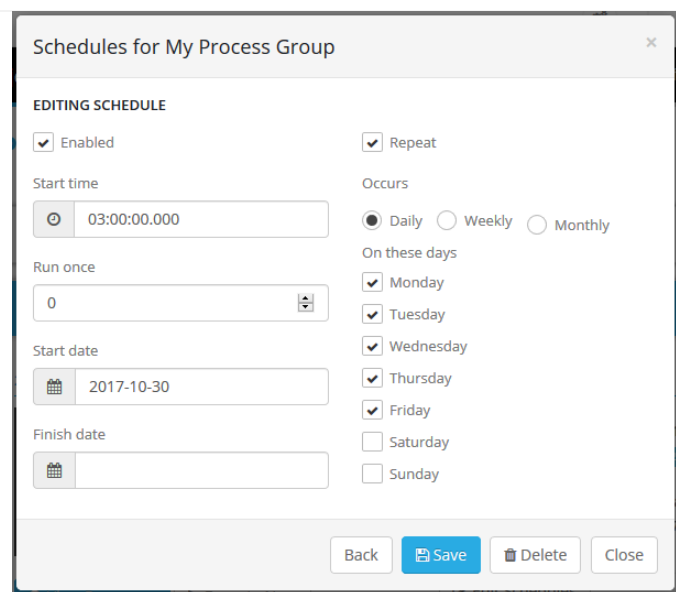


3. Here you can enable or disable the schedule, give it a start time, give it a duration (or only a start date with no end date), (optionally) set it to repeat, and on what basis to repeat.

Click 'Save' to create or update the schedule.

*or*

Click 'Delete' to delete the schedule.



*Schedules are based on the Time Zone of your Eightwire account – please ensure your Account's Time Zone is correct (see Edit Account).*

## Processes

### About Processes

A Process in Eightwire is the description of how to transfer data from one source table to one destination table.

Each Process contains information on which source columns map to which destination columns, additional calculated columns, behaviour settings, and a history of each time the Process has been executed.

In Eightwire, each individual execution of a Process is called a Batch. Or to be more accurate, a Batch is the execution of a specific task within a Process execution, so a single data transfer may actually be made up of a number of Batches, for example, read, transform, delete, update and insert. Each Batch contributes to the overall history or logging information for the Process execution.

While a Process is executing in Conductor, you do not need to stay logged in and are free to do other things. Your data will continue to be transferred on its own.

### Process Options

Conductor data transfer Processes are not rigid mechanisms. Instead they are smart, adaptive and self-healing. A Process is a definition of how you want a data transfer to work. When the actual data transfer takes place, Conductor uses the defined Process as a guide, but ultimately will attempt to make smart decisions about how to handle unexpected changes to data structures and data types.

The degree to which Conductor is allowed to exercise its own judgement is configurable. This allows for anything from a very rigid and predictable Process through to a very robust and heuristically driven Process.

The Process Options screen looks like this:

Options

Tolerate:

- ☒ Truncation
- ☒ Ambiguity
- ☒ Uncertainty
- ☒ Unusual
- ☒ Missing Source Columns
- ☒ Additional Source Columns

Column matching threshold

Medium

**Medium (45%) RECOMMENDED.** This is a good all-round setting, allowing a reasonable level of both accuracy and robustness. Lowest allowable matches include:

- Partial name match but non-guaranteed conversion
- Very similar name and guaranteed conversion
- Similar name and exact type match

### Truncation

When enabled, Conductor will allow oversize data to be shortened to fit into a smaller destination column.

### Ambiguity

When enabled, Conductor will allow ambiguous column mapping. For example, if a CSV contains two columns with the same name, Conductor will map them in the order it finds them. This may result in data arriving in the wrong destination columns, but the transfer will run without breaking.

### Uncertainty

When enabled, Conductor will allow conversion of one data type to another where the conversion is not always guaranteed to succeed. For example, converting text to date will succeed if the text column contains a date, but if it contains anything else it will fail. This may be required where the source data type cannot be determined or where the source database has been poorly designed.

### Unusual

When enabled, Conductor will allow some of the more unusual (but valid) data conversions. For example, number to date – which can legitimately be interpreted as a 'Julian date', but in most circumstances would be more likely to be an invalid conversion.

### Missing Source Columns

When enabled, Conductor will allow the data transfer to continue even if one or more of the source columns it expected was missing from the data repository. This can result in missing data at the destination, but in some situations this may be preferable to having the data transfer fail instead.

### Missing Destination Columns

When enabled, Conductor will allow a data transfer to continue even if one or more of the destination columns it expected was missing from the data repository.

*Note: Conductor will always tolerate additional columns in a data transfer. If new columns are added to both the source and destination repositories but Conductor has not been updated, this data transfer will detect the new columns, automatically map them (if they are compatible) and transfer the extra data.*

### Column Matching Threshold

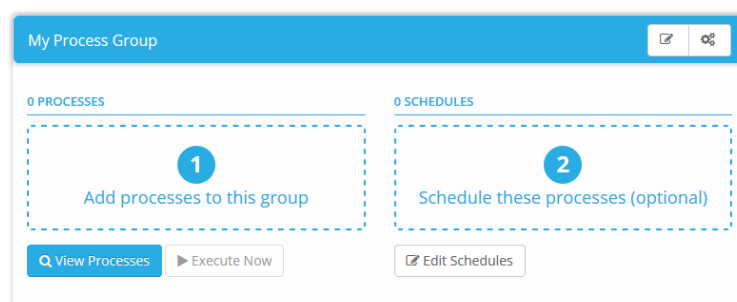
When Conductor first creates a Process for you, it automatically maps the source and destination columns for you. This 'smart mapping' uses a complex ranking system to calculate the best possible set of column mappings from the thousands of possible permutations. If the source and destination tables have identical columns and data types the mapping is easy. However, if the data types are slightly different Conductor will use the mappings that give the most accurate data conversions. If the column names are slightly different, Conductor will start using a progressive fuzzy-matching approach to determining the best set of column mappings. Eventually if the difference between

source and destination is too great, Conductor will not be able to come up with an acceptable mapping for these tables. The Column Mapping Threshold (acceptability threshold) is configurable – ranging from high (identical columns only will be mapped) to low (any compatible mapping will be considered). Generally a medium threshold gives a good balance between flexibility and accuracy.

This same smart mapping system is used when creating more than one Process at once – it determines the best possible destination table for each included source table.

## Tasks

To create, edit or delete Processes, you need to navigate to a Project, navigate to Process Groups, expand one Process Group panel on that page, then click 'View Processes'. You will be taken to the list of Processes for that Process Group (if any):



## Create New Processes

1.	Click 'Add Source' to select the source tables you would like to transfer.	<p><b>Processes in My Process Group</b></p> <p>Add source tables or queries</p> <p>+ Add Source</p>
2.	<p>This popup list shows you all of the source tables available to you from all of the source Data Stores in this Project.</p> <p>Select one or more tables from the list.</p> <p><i>To select multiple tables, hold the Ctrl or Shift key while you click.</i></p> <p>Click 'Add'.</p>	

3. Eightwire will compare the source tables you have selected against all the destination tables it knows about and suggest the best match for each table. If Eightwire can't find a good match it will instead suggest creating a new table on the destination.

(Optional) If you wish to change the suggested destination table or simply create the new table in a different Data Store, click on it and choose a new one.

If you are happy with all suggestions, click 'Accept All'. Alternatively, you can click the blue tick icon on each suggested Process.

*or*

To reject suggestions, either click 'Reject All' or click the red cross icon on each suggested Process.

**Processes in My Process Group**

There are no Processes in this group.

**Suggested Mappings** Accept All Reject All

dbo.FemaleFirstNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>
dbo.LastNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>
dbo.MaleFirstNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>

+ Add Source

↓

**Suggested Mappings** Accept All Reject All

dbo.FemaleFirstNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>
dbo.LastNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>
dbo.MaleFirstNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>

↓

**Suggested Mappings** Accept All Reject All

dbo.FemaleFirstNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>
dbo.LastNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>
dbo.MaleFirstNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>

*or*


**Suggested Mappings** Accept All Reject All

dbo.FemaleFirstNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>
dbo.LastNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>
dbo.MaleFirstNames	<input checked="" type="checkbox"/> <input type="checkbox"/>	New on Example Destination <span>New</span> <span>...</span>

## Edit a Process

1. To edit an existing Process, hover your mouse over the Process. You will see four buttons under your mouse cursor – click the left-hand 'edit' icon.

### Processes in My Process Group



2. Here, you will see two panels – the open Mapping panel and the closed Options panel underneath it.

*To learn how to modify columns, see the [Column Mapping](#) section of this manual.*

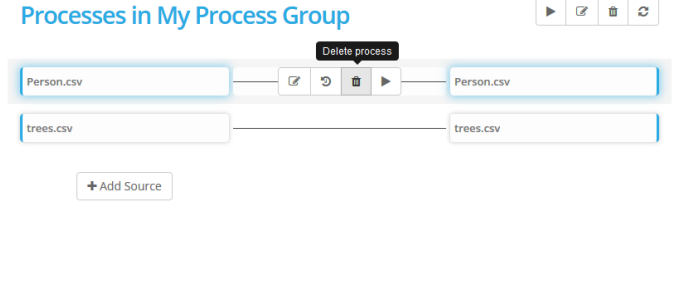
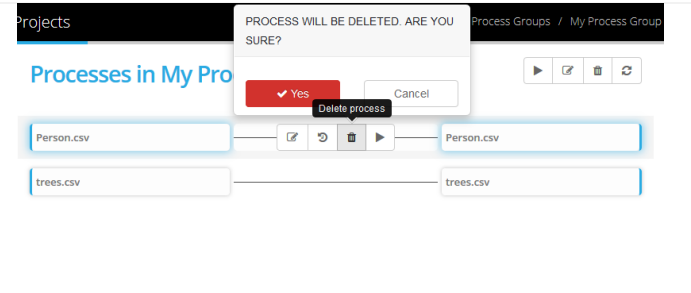
Click 'Options' to expand the Options panel.

### Editing Process Person.csv

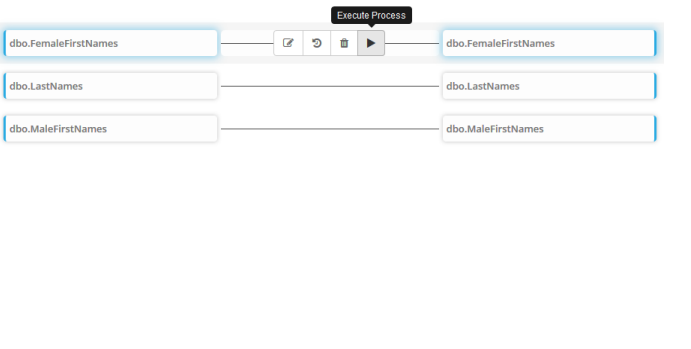
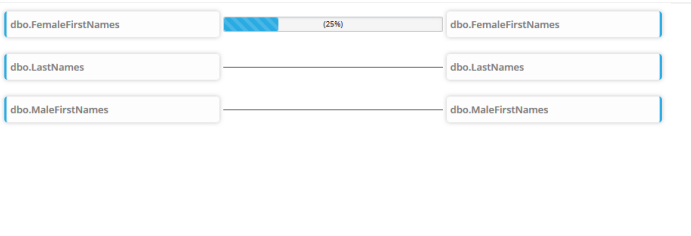
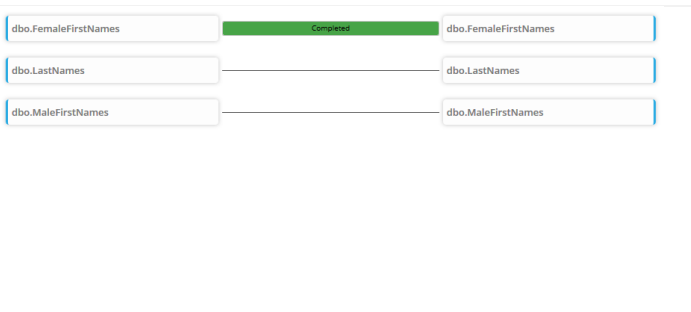
3. Change these options as required.  
  
Click 'Save' to update this Process.  
  
*or*  
  
Click 'Delete' to delete this Process.

### Editing Process Person.csv

## Delete a Process

1.	To delete an existing Process from the Process list, hover your mouse over the Process. You will see four buttons under your mouse cursor – click the 'delete' icon.	
2.	When prompted, click 'Yes' to confirm.	

## Execute a Process

1.	To execute an individual Process, hover your mouse over the Process in the Process list. You will see four buttons under your mouse cursor – click the right-most 'execute' icon.	
2.	The Process is now executing in the background. This page will show you the progress of the data transfer.	
3.	When the data transfer completes, you will see a green, orange or red indicator, showing success, success with warnings or failure respectively.	

## View Batch Activity

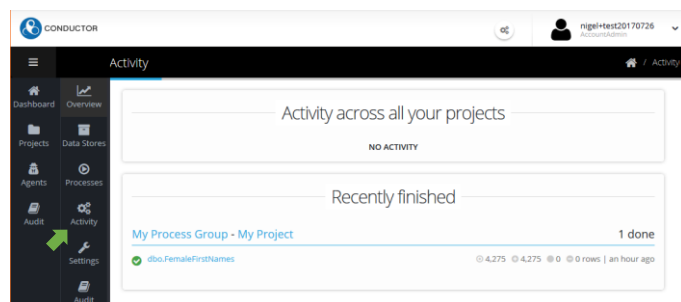
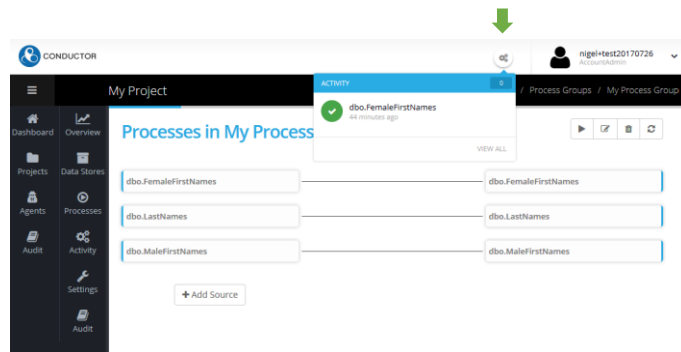
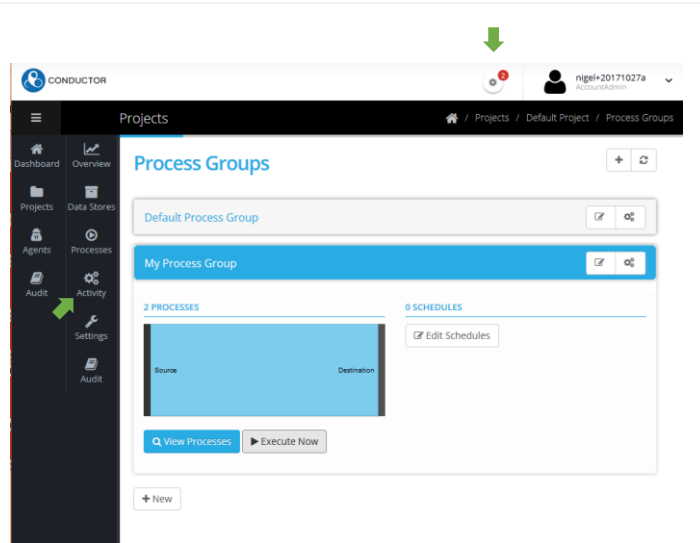
3. When Processes are executing, or have recently finished executing, you can check on their status from a number of places.

If you execute a single Process manually, you can monitor progress from that same page.

If you want to monitor all running Processes, there are two options:

While Processes are executing, the activity indicator at the top of the screen will animate. You can click on this icon to expand this mini-activity panel. You will see a selection of the most recently executed Processes and their status. *Not all Processes are shown here if there are a large number executing.*

The full activity page can be accessed by clicking 'Activity' in the left-hand navigation bar. This page will show you all current and recent activity and allow you to view the log for each Process execution.


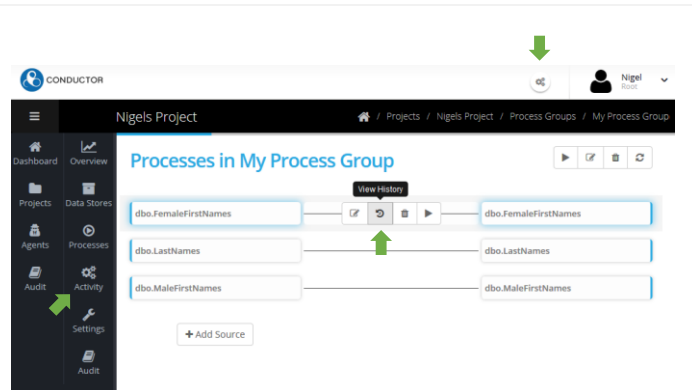




## View Batch History

1. In the *View Batch Activity* task above, we talked about three places where you can monitor activity. Each of these three places is an entry point for viewing Batch history.

For example, if you have executed a single Process, you can view the history of all executions of that Process by hovering your mouse over the Process and clicking the 'history' icon.



2. The popup window shows a list of all executions, their result, date and duration. It also shows which user ran it, or if it ran on a schedule.

In this example, the Process ran twice. The first time it completed successfully. The second time it completed successfully but there were warnings generated.

We can click on the top (second) execution to drill-down for more information.

History for dbo.FemaleFirstNames

Status	Date Started	Duration	Initiated By
⚠ Succeeded with warnings	01/08/2017 04:43 PM	2m 28s	nigel+test20170726
✅ Succeeded	01/08/2017 03:39 PM	2m 1s	nigel+test20170726

Refresh Close

3. This expanded view shows that this execution was actually made up of four batches – because the parent Process Group is set to 'Overwrite'. We see that the warnings came from the Transform Batch.

We click into the Transform Batch for more information.

History for dbo.FemaleFirstNames

Status	Date Started	Duration	Initiated By
▲ Succeeded with warnings	01/08/2017 04:43 PM	2m 28s	nigel+test20170726

Type	Status	Rows Involved	Batch ID
● Read	Succeeded	4275	1547757
▲ Transform	Succeeded with warnings	4275	1547758
● Delete	Succeeded	4275	1547759
● Insert	Succeeded	4275	1547760

● Succeeded	01/08/2017 03:39 PM	2m 1s	nigel+test20170726
-------------	---------------------	-------	--------------------

Refresh Close

4. Here we see individual entries in the log for this Batch, including the warning we were looking for.

*In this example, the table structure in our destination database was changed between the first execution and the second. Eightwire detected the unexpected change and handled it without any problems. However there was a risk of data truncation due to the nature of this particular change, so a warning was generated explaining it.*

History for dbo.FemaleFirstNames

Status	Date Started	Duration	Initiated By
▲ Succeeded with warnings	01/08/2017 04:43 PM	2m 28s	nigel+test20170726

Type	Status	Rows Involved	Batch ID
● Read	Succeeded	4275	1547757
▲ Transform	Succeeded with warnings	4275	1547758

Sev	Message
●	01/08/2017 04:43 PM Batch started by nigel+test20170726
▲	01/08/2017 04:44 PM 2 warnings found when comparing <b>Actual Source</b> to <b>Actual Destination</b> : The conversion of Name TEXT(50) to Name TEXT(20) is not guaranteed because the data types are not strictly compatible. The conversion of Name TEXT(50) to Name TEXT(20) could result in data loss through truncation.
●	01/08/2017 04:44 PM Batch succeeded

● Delete	Succeeded	4275	1547759
● Insert	Succeeded	4275	1547760

Refresh Close

## Column Mapping

### About Column Mapping

Each Data Store you define in Conductor contains a list of tables (found through *discovery*). Each table contains columns (found through *scanning*). Each Process defines how the data in one of those tables is moved from a source Data Store into a destination Data Store. To accomplish this, a Process needs to know how to map the data structures in the source table to the data structures in the destination table – this is called *Column Mapping*.

In most cases, the columns in the source and destination will be identical – usually because there is no existing destination table so one will be created with an identical structure to that of the source. However, in some cases the columns may be slightly different; in rare cases, substantially different.

When you create a Process, Conductor will attempt to pre-define the column mappings. If the columns on both sides are identical, they will all be mapped to their counterparts as expected. However, if the structures are different, Conductor will use *Smart Mapping* to determine the best mapping, while staying within the column mapping threshold set for each Process - see *Processes*.

### Smart Mapping

Many traditional data transfer tools and custom-written scripts tend to work on the principle that data structures cannot and should not change. Conductor works on the principle that data structures will change and often do in the real world. For this reason, Conductor does not use rigid column mapping. Instead Conductor uses *smart mapping*.

When the source and destination column structures are different, Conductor will try to construct a set of column mappings that best fits the data. As the difference between data structures gets greater, the accuracy of this matching is reduced. At some point the best possible accuracy achievable will fall below the *column matching threshold* defined for the Process and a mapping will not be possible.


If the column matching threshold is set at *Minimum*, Conductor will try everything possible to get a viable match, even seemingly unlikely matches. However, two columns must still be compatible before a mapping can exist. Mappings must not violate the tolerance settings for the Process.

For example, the following mapping has a high accuracy because only the case of the column names has changed. Any column matching threshold will allow this:

FullName	Text(100)	→	FULLNAME	Text(100)
DateOfBirth	Date	→	DATEOFBIRTH	Date
NumberOfCats	Number(10,0)	→	NUMBEROFCATS	Number(10,0)

The following mapping has a moderate level of accuracy because the column names are not quite identical but are subsets of each other and the types are not quite the same, but are convertible

without potential data loss. The fact that their columns are in a different order doesn't matter. This mapping would be possible with a *High* matching threshold or less:



The following mapping has a low accuracy because not only are the column names only similar at best, but some of the data types are only compatible through conversion, and are not guaranteed to always convert. For example, if the *DateOfBirth* column actually contains a textual date value then it will convert to a date at the destination, but if it contains some other text instead, it will fail. This mapping would be possible with a *Low* or *Minimum* matching threshold:



In the above example, the first two columns' names either phonetically sound similar or are subsets of each other and they have compatible or non-guaranteed convertible data types, so they can be mapped, but not with much accuracy. The rest of the columns are just too different and cannot be sensibly mapped, so are left unmapped by Conductor.

You can override Conductor's smart mapping to an extent by manually mapping two columns together. When you do this, Conductor will remap all of the other columns around them. However, this doesn't always guarantee that Conductor will respect these manually defined mappings – if the real-world data structures change enough to cause these two mappings to become incompatible or their mapping accuracy falls below the mapping threshold, Conductor will ignore your directions and attempt to do what it can with the data structures it has.

You can choose to not use one or more columns in a table. To do this, see the *Data Stores* section of this manual. Any columns marked 'do not use' will be excluded from smart mapping, even if they are a good match for a destination column. If you then re-allow these columns to be used, you will notice that the column mappings will adjust to accommodate them, and they will be mapped if they can be.

It is this behaviour that forms part of Conductor's self-healing feature set. The sort of changes you could expect to see in real life include renamed columns, new or removed columns and altered data types – all of these can usually be accommodated without human intervention.

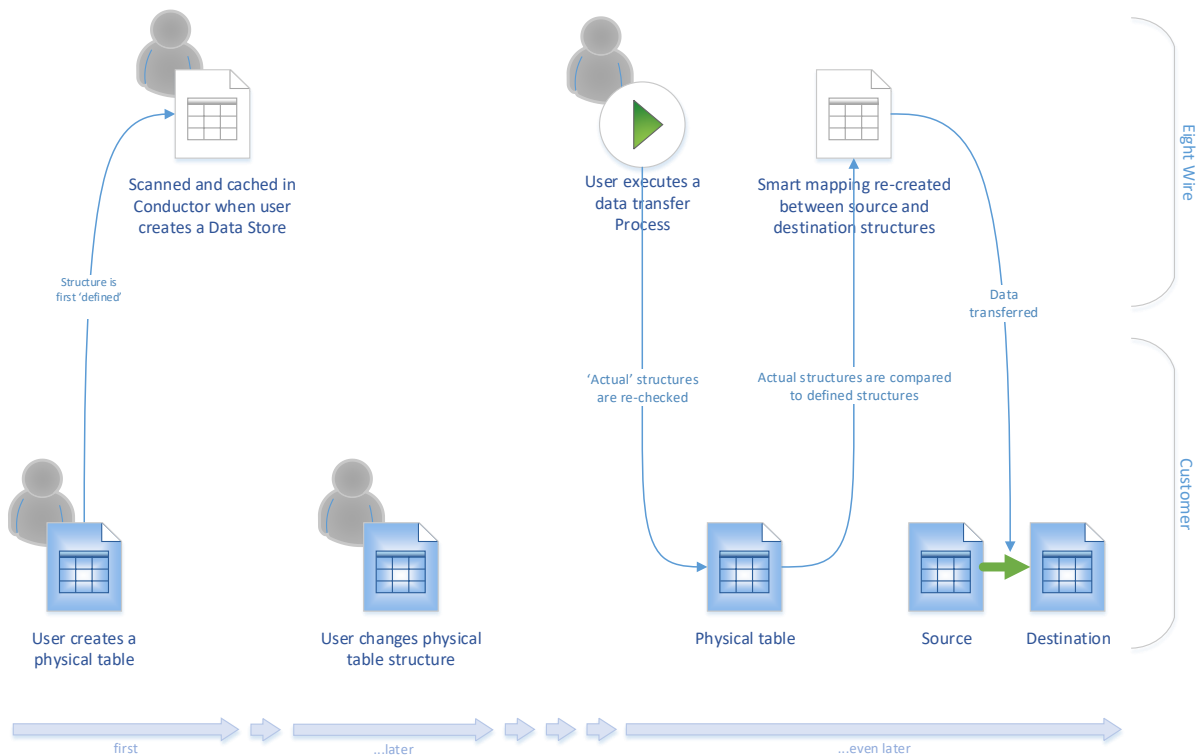
## Defined vs Actual Column Mapping

The following diagram shows a typical sequence of events over time.

First, a physical table is created in a database (or other type of platform). Then, a user logs onto the Eightwire portal and creates a Data Store. At that point Conductor scans the Data Store and stores its own copy of the table structure information – what columns and data types are used in the table.

Over time changes are made to that table, but Conductor may not be aware of them.

Later still, a user executes a data transfer Process using this Data Store (either as a source or as a destination). At this time, Conductor re-checks the physical table to see if the *actual* structure is the still the same as the previously *defined* structure. It then creates a new smart mapping, based on the actual source and destination structures and the defined source and destination structures – making sure it takes into account any preferential mappings the user has created, but also adjusting for any changes in the structures over time. The resulting mapping is what is actually used to perform the data transfer process.



This diagram does not show a 'Rescan' action, but ideally one should have occurred just after the user changed the physical table structure.

If you have configured your Conductor Data Store to allow automatic re-scanning, then Conductor will periodically rescan your tables and adjust its own copy of the *defined* structure information. It's important to keep this information up-to-date because this is what is used when you are creating and editing Processes – if this information is out-of-date then the Process you think you have built might actually look different when you execute it.

You can manually re-scan a Data Store at any time.

## Tasks

To edit column mappings, you need to navigate to a Project, navigate to and expand one Process Group, click 'View Processes', then hover your mouse over one Process and click on the 'Edit' icon that appears on the line between the source and destination for that Process. You will be taken to a page showing both the column mappings and options for that Process. This page looks similar to this:

### Editing Process dbo.LastNames

## Change column mapping

1. To override the automatically generated 'smart mapping' for a column, first click on a column on the left-hand side.


2. Next, click the corresponding column on the right-hand side. This will create a *preferred mapping* between these columns.

*Note that all other columns have also been automatically remapped to account for your new preferred mappings.*

3. The circular arrow icon between the source (left) and destination (right) columns indicates a preferred mapping.

(Optional) To remove the 'preferred' status from this mapping, click the *preferred mapping icon*.

When you are finished, **click 'Save'**.



### Prevent column mapping

See the task *Browse and Edit Data Structures* in the *Data Stores* section of this manual for information on how to mark individual columns as 'do not use'.

Columns that are marked 'do not use' are shown on the mapping page, but cannot be mapped.

## Column Manipulation

### About Column Manipulation

In addition to simply moving data, you may wish to manipulate that data during the data transfer. There are a number of ways to do this, including either filtering out unwanted rows of data, or adding additional columns of data.

### Filtering

When used, one or more filter conditions on a table serve to remove any rows of data that do not meet **all** of the filter conditions. Filter conditions are attached to individual columns but are applied collectively to the whole table.

For example, the filter condition 'Age > 20', applied to the *Age* column on a table, will mean that only rows that have a value in the *Age* column greater than 20 will be sent to the destination – all other rows will simply be discarded. If we then add the additional filter condition 'Name = "Sarah"' to the *Name* column, this now only sends rows to the destination that have Age greater than 20 **and** the Name "Sarah". If no filters are applied, all rows will pass through the filter without any being removed.

Filters can be applied in three different places – during the Read phase, during the Transform phase, and during the Write phase. These all have the same effect, and 'stacking' them will progressively remove more rows as each filter is applied. This section of the manual describes Filters applied during the Transform phase. For information on how to apply Filters in the Read or Write phases, see the *Manage Data Store Filters* task in the *Data Stores* section of this manual.

### Constants

A *Constant* is a fixed data value - it is not calculated and it does not change. It may be desirable to simply add a number or a piece of text to each row of data being transferred. For example, when transferring the type of data into one destination from multiple sources, it may be useful to include the name of the source as a Constant column for each respective Process. When the data arrives at its destination each row will have one additional column containing a hard-wired value.

Date	Temperature		Wellington
10/11/2017	18		Wellington
11/11/2017	17		Wellington
12/11/2017	17		Wellington
13/11/2017	15		Wellington
14/11/2017	14		Wellington
15/11/2017	15		Wellington

### System Metadata

System Metadata is similar to a Constant, except that instead of specifying an arbitrary constant data value, we will use a data value representing information generated by Conductor relating to the



data transfer. There are many different metrics available, describing different aspects of a data transfer and its related information. For example, adding the *BatchExecuteDate* System Metadata column to the Process will add the data transfer date to every row processed. System Metadata values are calculated at the beginning of each Process execution.

Date	Temperature		BatchExecuteDate
10/11/2017	18	+	16/11/2017 4:08pm
11/11/2017	17		16/11/2017 4:08pm
12/11/2017	17		16/11/2017 4:08pm
13/11/2017	15		16/11/2017 4:08pm
14/11/2017	14		16/11/2017 4:08pm
15/11/2017	15		16/11/2017 4:08pm

## Expressions

An Expression is a mathematical formula, used to calculate the data value that should appear in each row of the Expression column. Expressions can be used with text, numbers, dates and other data types. Expressions operate similarly to spreadsheet formulas – they can contain mathematical operators, functions and conditions. Expressions in Conductor have their own syntax, but for anyone familiar with spreadsheet formulas, they should be simple to learn.

Expressions are a large topic in themselves, so this manual will not cover them in detail. Please instead download or request a copy of the *Expressions Guide*.

Date	Temperature		"Hello " + "World"
10/11/2017	18	+	Hello World
11/11/2017	17		Hello World
12/11/2017	17		Hello World
13/11/2017	15		Hello World
14/11/2017	14		Hello World
15/11/2017	15		Hello World

## Tasks

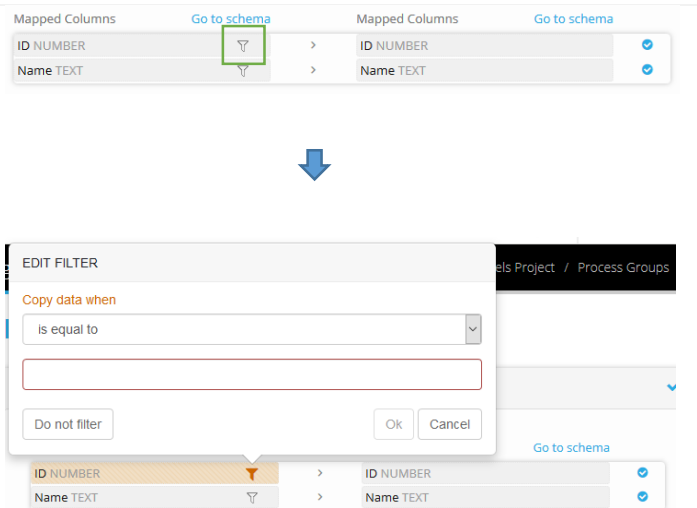
Columns are manipulated from the same page as column mappings. To manipulate columns, you need to navigate to a Project, navigate to and expand one Process Group, click 'View Processes', then hover your mouse over one Process and click on the 'Edit' icon that appears on the line between the source and destination for that Process. You will be taken to a page showing both the column mappings and options for that Process.

### Filtering a Column

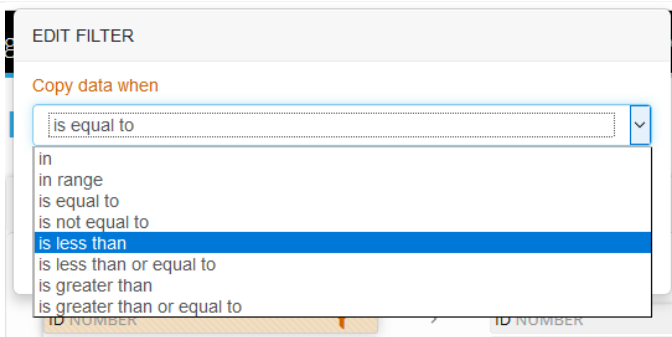
1. To add a filter criteria, choose the column it relates to and click the filter icon on the source (left hand side) column.

You will see a pop-up window, asking you to provide filter details.

*Rows that **do not** meet this criterion will be discarded during the data transfer.*

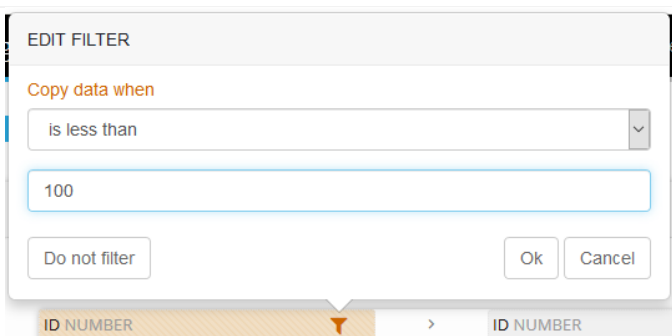


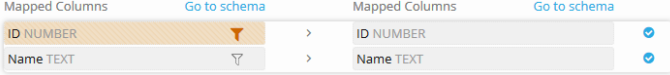
2. In the first box, select the condition.



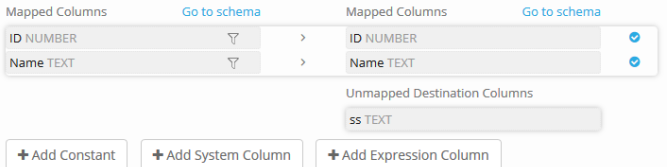
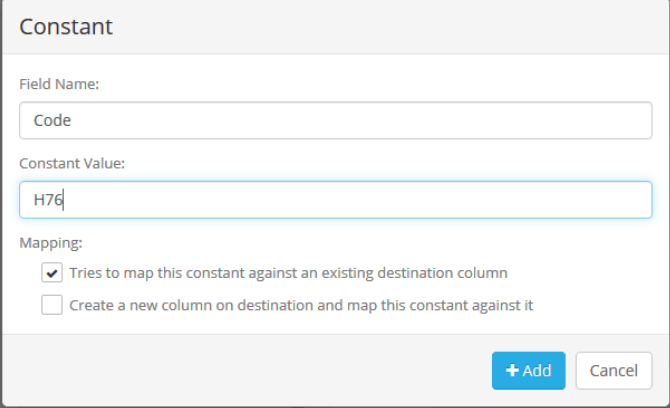
3. In the second box, select the matching value.

*In this example, we want to include rows where the ID field contains a value **less than 100**. This means that rows with ID values greater than or equal to 100 will be discarded.*



<p>4.</p>	<p>Click 'OK' to add the column filter.</p> <p>To edit the same filter, click the filter icon again.</p> <p>When you are finished, <b>click 'Save'.</b></p>	 <p><i>Unsaved changes appear in a different colour.</i></p> <p><i>A dark filter icon indicates a filter is present.</i></p>
-----------	---	--

### Adding a Constant

<p>1.</p>	<p>To add a Constant column, click 'Add Constant'.</p>	
<p>2.</p>	<p>In the popup window that appears, enter a name for the column and the constant value you wish to add to every row in this new column.</p> <p><i>You may see two additional checkboxes on this popup – use these to either automatically map this column to a destination column or create a new matching destination column. These options are not available if the destination table already exists.</i></p> <p>Click 'Add' to add this column.</p>	

3. Your new column appears on the left-hand side of the mapping.

See the *Change Column Mapping* task in the previous section of this manual to map your new column to a destination column.

To edit the name or value of the Constant column, click the *Pi* ( $\pi$ ) icon on the column.

### Adding System Metadata

1. To add a System Metadata column, click 'Add System Column'.

2. In the popup window that appears, enter a name for the column.

Select the type of system metadata you wish to add to every row in this new column.


*You may see two additional checkboxes on this popup – use these to either automatically map this column to a destination column or create a new matching destination column. These options are not available if the destination table already exists.*

Click 'Add' to add this column.

3. Your new column appears on the left-hand side of the mapping.


See the *Change Column Mapping* task in the previous section of this manual to map your new column to a destination column.

To edit the column's name or type of system metadata, click the gears icon on the column.



### Adding an Expression

1. To add an Expression column, click 'Add Expression Column'.



2. In the popup window that appears, enter a name for the column.

Enter the Expression formula.

See the *Expressions Guide document for information about using Expressions.*

You may see two additional checkboxes on this popup – use these to either automatically map this column to a destination column or create a new matching destination column. These options are not available if the destination table already exists.

Click 'Add' to add this column.

Expression

Field Name:

Message

Expression:

"Hello " + "World"

Mapping:

☒ Tries to map this expression column against an existing destination column

☐ Create a new column on destination and map this expression column against it

+ Add


Cancel

3. *If there is a problem with your Expression formula, you will need to correct it before progressing.*

Your new column appears on the left-hand side of the mapping.

See the *Change Column Mapping* task in the previous section of this manual to map your new column to a destination column.

To edit the columns name or formula, click the *expression* icon on the column.



# Reporting

## About Reporting

Conductor includes a number of pre-built reports. Some of these reports are specific to your account, and some apply to multiple accounts based on any Account Groups you are participating in.

Users with the Account Administrator or Account Viewer role will be able to access the Reports section of Conductor.

## Account Groups

If your organisation is part of a larger community, such as a government sector or federated group of companies, with your permission, your Account may be part of an Account Group which includes the other organisations you are collaborating with.

A number of reports operate over an entire Account Group, allowing a high-level view of what is happening across the entire group rather than just within one Account. Not every participant within a group will have access to all of the reports that cover that group – if you are part of a group and wish to have access to your group's reports, please contact the Eight Wire Support team.

## The Reports

The following reports are available to authorised users and accounts. This description of each report relates to its purpose and how its content is generated.



### Aged Users Report

Use this report to highlight user accounts that are no longer needed and should be removed.

This report shows only those users who have either never signed in or not signed in within the past 90 days, as at the date the report was run.

All dates and times are localised to your time zone, as set in the Eight Wire settings for your account.

A user account will show as locked out if someone has used the wrong password five times in a row. Locked accounts can be unlocked using the password reset procedure. Accounts that remain locked or are frequently locked could indicate unauthorised attempts to guess a password. This type of activity should be investigated.

Unused user accounts could pose a security risk to your organisation. If these accounts belong to people no longer employed or contracted to you, these people should no longer have access to your data structures and business processes. You should consider removing these users if they are no longer needed.



## Account Summary Report

Use this report to determine the security posture and general usage of all Eight Wire accounts that belong to the same Eight Wire account groups as you.

Accounts are listed in this report at the request of the nominating organisation or at each accounts own request. Individually listed accounts will only have access to this report if the nominating organisation has additionally requested this.

All information shown in this report is current at the time the report was run, except 'Last Data Transfer' date, which is calculated hourly. Accounts are only listed if they are current at the time the report was run.

'Two Factor Authentication' indicates whether an account has enabled this additional security feature for all of its users. Eight Wire recommends that all accounts should use Two Factor Authentication to help prevent unauthorised access.

The number of users, projects and agents gives an indication of the scale of activities for each listed account.

All dates and times are localised to your time zone, as set in the Eight Wire settings for your account.

'Last Login' shows the date and time of the last user sign in for each listed account. Accounts with infrequent user activity may indicate lack of use, or that everything is running on a schedule and does not require user intervention.

'Last Data Transfer' shows the date and time of the last data transfer executed, either manually or on a schedule, either internal or between accounts (regardless of an accounts presence in this report). For transfers between accounts, it doesn't matter which account owns the process, this activity will still appear here if it is the latest transfer relating each listed account. This metric simply indicates if any data transfers have taken place recently or not.





## Data Movement Report

Use this report to see all data sharing agreements and data transfers between any Eight Wire accounts that belong to the same Eight Wire account groups as you.

You may choose the reporting period. The reporting period starts at midnight on the first date (inclusive) and runs through to 11:59:59pm on the last day (inclusive). All dates and times are localised to your time zone, as set in the Eight Wire settings for your account.

Accounts are listed in this report at the request of the nominating organisation or at each accounts own request. Individually listed accounts will only have access to this report if the nominating organisation has additionally requested this.

Accounts are only listed if they are current at the time the report was run and have been included in the reporting group by Eight Wire. Deleted accounts will not be shown, even if they existed during the reporting period or were previously in the reporting group but have subsequently been removed.

All information shown in this report is aggregated and valid for the reporting period only, except 'Tables', which is current at the time the report was run.

'Tables' shows the total number of tables which have inter-account data transfer processes defined, regardless of whether these processes have ever been executed.

'Transfers' shows the total number of discrete table-level transfers (batches) that took place, within the reporting period, between the listed account and any other account included in this report, across all tables and data transfers for each sharing agreement. Internal (intra-account) transfers are excluded. Inter-account transfers between a listed account and an account not included in this report are also excluded.

'Rows' shows the total number of discrete rows that were transferred across all 'Transfers' as described above.

If a data sharing agreement has no data transfers within the reporting period, this will show as 'no data transferred'. Unused data sharing agreements should be paused or permanently removed to limit unnecessary access to or from external organisations.



## Data Security Report

Use this report to review security and privacy related information for all data sharing agreements and data transfers between any Eight Wire accounts that belong to the same Eight Wire account groups as you.

All information shown in this report is current at the time the report was run, except 'Last Data Transfer' date, which is calculated hourly.

Accounts are listed in this report at the request of the nominating organisation or at each accounts own request. Individually listed accounts will only have access to this report if the nominating organisation has additionally requested this.

Accounts are only listed if they are current at the time the report was run.

All dates and times are localised to your time zone, as set in the Eight Wire settings for your account.

Each report section shows information relating to one combination of source and destination Data Store, where Processes exist to transfer data (regardless of whether they have ever been executed or not), where these transfers are defined between accounts (not internal to one account) and are defined using a data sharing agreement.

For each report section, the account name, project name and data store name are shown for both the sending and receiving data stores.

'Personally Identifying and Notifiable Columns' shows any source columns that have been directly or indirectly tagged as PII (Personally Identifying Information) or as Notifiable by the sending account. This provides an indication of the sort of information that is potentially being shared that may be of interest from a privacy or security perspective.

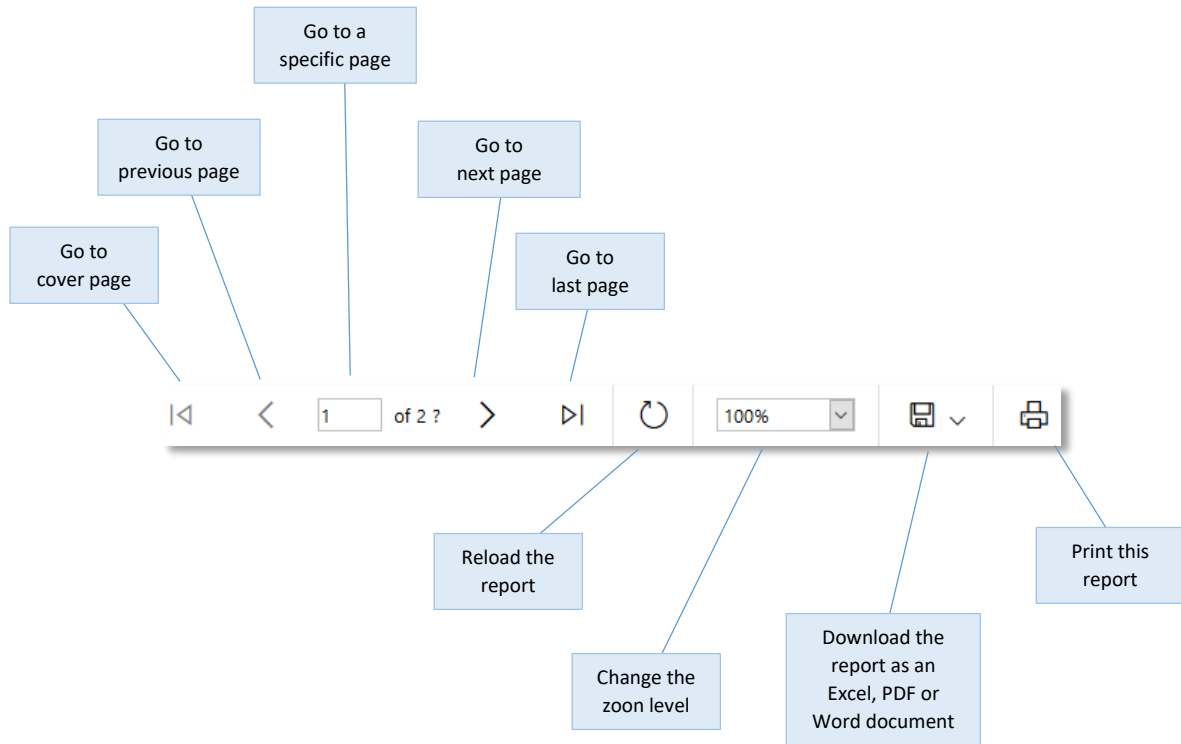
'Data Sovereignty' indicates whether any agreed data sovereignty exists. If a region is listed, any data transfers will be restricted to within this region. If no region is listed, data may move through any Eight Wire processing region as system load dictates. 'Mismatch' indicates that agreement on data sovereignty has not been reached and Eight Wire will prevent any data transfers until agreement has been reached.

'Sharing Agreement' indicates whether this data sharing agreement is active or paused. Any text entered for this data sharing agreement is shown here – this would normally be either a reference to a legal agreement or the text of the agreement itself. If no agreement text has been entered, this may indicate poor practice or that data is being shared without a legal agreement to do so.

'Last Data Transfer' will show either the date of the last data transfer executed between this pair of Data Stores, or it will show 'No transfers within the last 90 days' where no transfers have taken place within the past 90 days, from the date the report was run. Unused inter-account data transfer processes should be removed and unused data sharing agreements paused or permanently removed to limit unnecessary access to or from external organisations.

## Using Reports

Each report provides different information, but each report is navigated, downloaded and printed in the same way. Whenever you are viewing a report within the web portal, you will see the same toolbar at the top of the reports panel on the page:



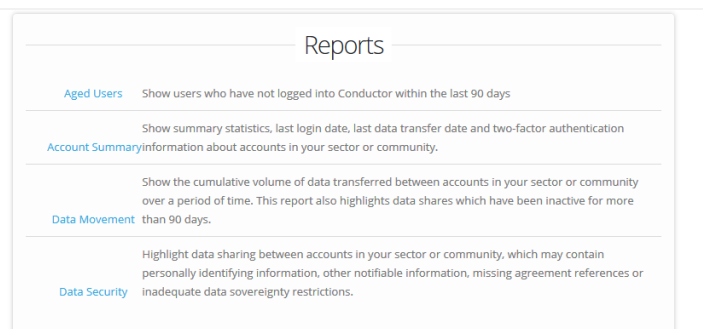
## Tasks

To view reports, you need to be on the **Account Dashboard** page. From the Dashboard, click on the 'Reports' icon in the Navigation bar.

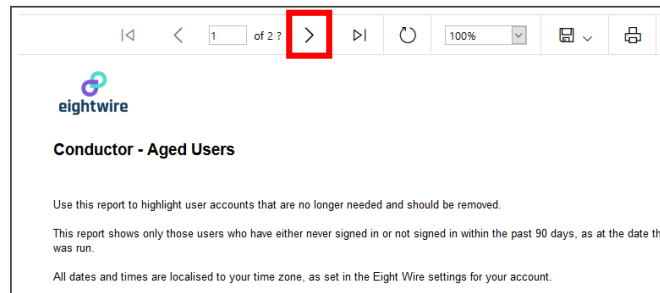
### View a Report in the Web Browser

1. Click the name of the report you wish to view.

You will be taken to that report.

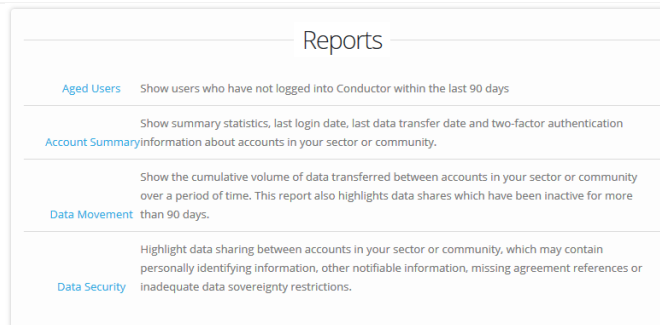


2. *The cover page provides important information about the report and how to interpret it.*
- Click the 'right arrow' icon at the top of the report to move from page to page.

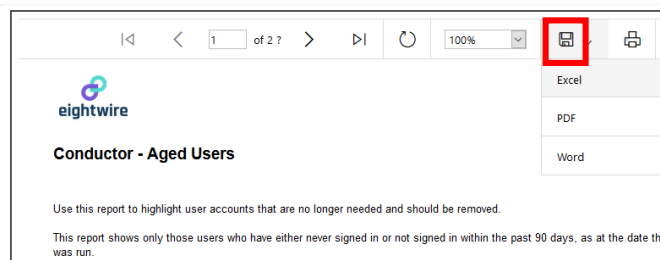


## Download a Report in Excel, PDF or Word Format

1. Click the name of the report you wish to view.
- You will be taken to that report.



2. Click the 'download' icon at the top of the report to show your download options.
- Click the file format you want and your download will begin.



## Audit Log

### About Audit Logging

Every time a user changes something within Conductor, the change is recorded. This provides an immutable record of what was changed, when and by whom. For example, if you login to the web portal, open a Project and change its description, an entry is written to the audit log, showing what the Project record looked like before and after the change, your user name (because you made the change) and the time and date of the change. Audit logging occurs when you change any object within Conductor.

The Audit Log can be viewed by Account Administrators and Project Administrators, by clicking the 'Audit Log' icon on the navigation bar or on the dashboard. The Audit Log page allows you to filter the audit log by any of several criteria to show only the information you want. The audit log page looks like this:

### Audit Logger

Project: Nigels Project User: Nigel From: 01 / 11 / 2016 To: 30 / 11 / 2016 Entity: Process Audit action: All Entity ID:  Filter

Date	Entity	Entity ID	Action	User
08/11/2017 01:54 PM	Process	9248	Insert	Nigel
08/11/2017 01:54 PM	Process	9248	Update	Nigel
08/11/2017 01:55 PM	Process	9249	Insert	Nigel
08/11/2017 01:55 PM	Process	9249	Update	Nigel
08/11/2017 01:55 PM	Process	9250	Insert	Nigel
08/11/2017 01:55 PM	Process	9250	Update	Nigel
08/11/2017 03:44 PM	Process	9260	Insert	Nigel
08/11/2017 03:44 PM	Process	9260	Update	Nigel


State Before

```
<Process ID="9248" Name="TEST-CONNECTION for DataStore #2576"
DateCreated="2017-11-08T00:54:17.323" DateModified="2017-11-08T00:54:17.323"
Description="TEST-CONNECTION Process" AccountID="2" Command="" IsPartial="1"
MappingThreshold="1" TolerateTruncationConversion="1"
TolerateAmbiguousConversion="1" TolerateUncertainConversion="1"
TolerateUnusualConversion="1" TolerateMissingEndPointColumns="1"
```

State After

```
<Process ID="9248" Name="TEST-CONNECTION for DataStore #2576"
DateCreated="2017-11-08T00:54:17.323" DateModified="2017-11-08T00:54:17.350"
Description="TEST-CONNECTION Process" AccountID="2" Command="" IsPartial="1"
MappingThreshold="1" TolerateTruncationConversion="1"
TolerateAmbiguousConversion="1" TolerateUncertainConversion="1"
TolerateUnusualConversion="1" TolerateMissingEndPointColumns="1"
```

You may also export the filtered audit log records as a CSV or PDF document. To do this, first define the scope of the search, click filter, and then click the download icon on the top-right of the table, as shown below:



The length of time audit logs are kept by Eightwire is defined by either legislation or customer contract. Please contact Eightwire Support for more details.

## Appendix A - Connectors

The Eightwire Conductor platform is designed to read and write data to and from a wide variety of different data platforms, from files to databases. When you create a Data Store in Conductor, you will be asked how your data is stored:

How is your data stored?

**APPLICATION**

☒ Salesforce

**FILE**

☐ Delimited Text      ☒ Delimited Text Dropfile      ☐ Excel Spreadsheet      ☒ Excel Spreadsheet Dropfile

☐ Fixed Width      ☐ Folder      ☐ FTP Delimited Text      ☐ FTP Excel Spreadsheet

☒ Hadoop Distributed File System

**NOSQL**

☒ Cloudant      ☐ CouchDB      ☒ DocumentDB      ☒ MongoDB

**RELATIONAL DATABASE**

☐ ACCESS      ☐ ANSI SQL      ☐ DB2 Generic      ☐ DB2 ISeries

☒ Hadoop HIVE      ☐ Interbase      ☐ MYOB      ☐ MySQL

☐ Netezza      ☐ Oracle      ☐ PostgreSQL      ☒ Redshift

☐ SQL Server      ☐ SYBASE

Each Data Store makes use of a *Connector* to access data on your chosen platform. Each Connector is required to perform the same basic tasks but they accomplish those tasks in quite different ways. For example, when asked to write data, an Excel Spreadsheet Connector will create an Excel file containing the data and store it in a folder system, whereas an SQL Server Connector will execute database commands to write the same data into the database. From a user perspective, both achieve the same task, but internally they work quite differently.

Because all data formats have differing capabilities and limitations, so do the Conductor connectors that work with them. Where possible Conductor will attempt to bridge the capability gap by inferring the missing information from the data, even where this is not natively provided by the data format itself.

For example, the Delimited Text format does not provide any way to designate a primary key column. Instead, Conductor analyses some of your data and infers an appropriate column for you. Conversely, most database platforms explicitly allow you to specify a primary key, which Conductor will detect and use.

## Destination Actions

The key tasks a Connector performs are reading and writing. Reading is usually straightforward, but writing can be complex, depending on the platform and on the other data being written at the same time.

**Append** – Add new data onto the end of any existing data.

**Overwrite** – Remove any existing data, then add the new data.

**Partial Merge** – Given knowledge of the unique keys in both the source and destination tables, add (insert) new data and update existing data. Any other destination rows are untouched.

**Full Merge** - Given knowledge of the unique keys in both the source and destination tables, add (insert) new data, update existing data and delete any existing destination data that is not also in the source. This is recommended for database tables with referential integrity constraints.

**Incremental** – The maximum numeric, date or datetime value for a nominated column in the destination table is compared with each source row. Any source rows whose values in that column are greater than the maximum destination value are appended to the destination table. This is recommended for tables that are added to, but never updated.

For each Connector, the following *Destination Actions* are available:

	Read	Append	Overwrite	Partial Merge	Full Merge	Inc. Load
<b>Application</b>						
SalesForce	✓			✓		
<b>File</b>						
Delimited Text	✓		✓			
Delimited Text Dropfile	✓					
Excel Spreadsheet	✓		✓			
Excel Spreadsheet Dropfile	✓					
Fixed Width	✓		✓			
FTP Delimited Text	✓		✓			
FTP Excel Spreadsheet	✓		✓			
HDFS (Delimited Text)	✓		✓			
<b>NoSQL</b>						
Cloudant	✓	✓	✓			
Couch DB	✓	✓	✓			
Document DB	✓	✓	✓			
Dynamo DB *	✓	✓	✓			
Mongo DB	✓	✓	✓			
<b>Relational Database</b>						
ACCESS	✓	✓	✓	✓	✓	✓
ANSI SQL	✓	✓	✓	✓	✓	✓
DB2 Generic	✓	✓	✓	✓	✓	✓
DB2 iSeries	✓	✓	✓	✓	✓	✓
Hadoop HIVE	✓	✓				
Interbase	✓	✓	✓	✓	✓	✓
MYOB	✓					
MySQL	✓	✓	✓	✓	✓	✓
Netezza *	✓	✓	✓	✓	✓	✓
Oracle	✓	✓	✓	✓	✓	✓
PostgreSQL	✓	✓	✓	✓	✓	✓
Redshift	✓	✓	✓	✓	✓	✓
SQL Server	✓	✓	✓	✓	✓	✓
SYBASE	✓	✓	✓	✓	✓	✓
<b>Binary</b>						
Folder (file transfer) **	✓		✓			

\* Only available in the Beta Programme

\*\* Only available upon request



## List of Connectors

This section contains Connector-specific notes and instructions.

### Application Connectors

Generally, Conductor deals with databases, files and other data storage platforms. However, this special class of connectors is designed to provide an abstraction layer over application-specific data, so that it can be managed like any other data platform.

#### *SalesForce.com*

Uses the SalesForce.com API and is available as a direct connection only (not via Agent). Table structures cannot be modified. Does not support deletion of records – they should be flagged as deleted as per SalesForce.com advice. The use of this connector may be subject to various throughput limitations imposed by SalesForce.com – see the SalesForce.com website for documentation on API limitations.

### File Connectors

The *path* property for this Connector should be the path to the folder containing the files, **not** a path to a single file. If used with an Agent, the path should be relative to where that Agent is installed. Supports UNC paths, provided the correct network trusts and permissions are in place.

#### *Delimited Text*

Handles delimited text files in a number of formats, including Comma Delimited Text (CSV) and Tab Delimited Text (TSV). The connection properties allow you to specify which row (if any) the header appears in, which row the data starts in, what the column and row delimiters are and which character to use as a text-qualifier. Supports the Wildcard feature.

This connector can be used with files stored on disk, network file store, or FTP/SFTP server.

Disk – the Agent reads or writes files in the specified *path*, using its own security credentials.

FTP – use a username and password to read or write files on an FTP server, in the specified path.

SFTP – use a username and password or a username and public key file (a .pfx file stored with the Agent) to read or write files on an SFTP server (FTP over SSH, using OpenSSH), in the specified path.

#### *Delimited Text Drop File*

**Supports read operations only.** Provides a CSV file drag and drop facility for source Data Stores. This functions exactly like the Delimited Text connector, but Processes built on this Data Store are initiated by manually dropping files onto an area of the web portal, they cannot be scheduled or executed by API call.

When creating this Data Store, select 'In the cloud' when asked – you will not be using an Agent even though your files are behind a firewall.

#### *Excel Spreadsheet*

Handles Excel (XLSX) files. The connection properties allow you to specify which row (if any) the header appears in and which row the data starts in. This connector cannot be used to handle older XLS file types. **Do not use this connector to handle CSV or TSV files.**

This connector can be used with files stored on disk, network file store, or FTP/SFTP server.

Disk – the Agent reads or writes files in the specified *path*, using its own security credentials.

FTP – use a username and password to read or write files on an FTP server, in the specified path.

SFTP – use a username and password or a username and public key file (a .pfx file stored with the Agent) to read or write files on an SFTP server (FTP over SSH, using OpenSSH), in the specified path.

#### *Excel Spreadsheet Drop File*

**Supports read operations only.** Provides an Excel Spreadsheet (.xlsx) file drag and drop facility for source Data Stores. This functions exactly like the Excel Spreadsheet connector, but Processes built on this Data Store are initiated by manually dropping files onto an area of the web portal, they cannot be scheduled or executed by API call.

When creating this Data Store, select 'In the cloud' when asked – you will not be using an Agent even though your files are behind a firewall.

#### *Fixed Width*

Handles fixed-width text files. When a Data Store of this type is created, Conductor attempts to infer the positions of the columns as well as their data types. Some Fixed Width files may not detect correctly if their format is not obvious – in this case, it is possible to adjust the columns, start positions and lengths in the *Browse Data Store* page. The connection properties allow you to specify which row (if any) the header appears in, which row the data starts in and which character to use as padding. Supports the Wildcard feature.

#### *Hadoop Distributed File System (HDFS) for Delimited Text*

Allows access to delimited text files stored on HDFS. Available as a direct cloud-to-cloud connection only.

### NoSQL Database Connectors

While Conductor has support for a number of NoSQL data platforms, it should only be used where the NoSQL platform stores data in a pseudo-tabular way. Each document should equate to a single row of data, with data stored as key-value pairs without 'nesting'. These restrictions will be addressed in an upcoming version of Conductor.

#### *Cloudant*

Supports direct connections only. Requires *Database*, *Username*, *Password* and *EntityColumnName* connection properties.

#### *CouchDB*

Supports both direct and Agent-based connections. Requires *Database*, *Username*, *Password*, *Port*, *URL* and *EntityColumnName* connection properties.

#### *DocumentDB*

Supports direct connections only. Requires *URL*, *Access Key* and *Database* connection properties.

#### *Dynamo DB*

Supports direct connections only. Only available to participants in the Beta Programme. Requires *AWS Region*, *Access Key* and *Secret Key* connection properties.

#### *MongoDB*

Supports direct connections only. Requires *URL*, *Database*, *Username* and *Password* connection properties.

## Relational Database Connectors

Almost all relational database Connectors provide the full range of feature support, including full referential integrity detection and compliance during multi-table data transfers. Some of the Connectors in this group may seem out of place, such as HIVE or MYOB, but they are grouped here because these connectors operate much like a relational database in the way they access data.

### *Access*

Supports ODBC and OleDb via an Agent with appropriate drivers installed.

### *ANSI SQL*

If you have a database which is not directly supported by Conductor, but is 100% ANSI compliant, you could try this connector. It is designed to issue ANSI compatible commands against any SQL-based database system, via an Agent with the appropriate ODBC or OleDb drivers installed. **Caution is advised** – while every effort has been made to ensure the correct functioning of this Connector, you will ultimately be using it in an untested configuration. It is also common practice for most database manufacturers to claim their products are ANSI compliant, but in Eightwire's experience almost none are truly 100% compatible.

### *DB2 Generic*

The standard DB2 Connector. Supports ODBC and OleDb via an Agent with appropriate drivers installed.

### *DB2 iSeries*

The iSeries variant of the DB2 Connector. Supports ODBC and OleDb via an Agent with appropriate drivers installed.

### *Hadoop HIVE*

The HIVE database connector for Hadoop. Available as a direct cloud-to-cloud connection only.

### *Interbase*

Supports ODBC via an Agent with appropriate drivers installed.

### *MYOB*

**Supports read operations only.** Uses the MYOB ODBC-compliant 32-bit driver, via a 32-bit Agent, to access MYOB application data. This driver has a number of limitations, so some care and patience is required. Avoid transferring data to or from multiple tables at the same time if possible – this can cause timeouts and table locks. The Agent must be installed on the MYOB server. Does not require any connection properties.

### *MySQL*

Supports ODBC and OleDb via an Agent with appropriate drivers installed.

### *Netezza*

Supports ODBC and OleDb via an Agent with appropriate drivers installed. This Connector is only available to Beta Programme participants.

### *Oracle*

Supports Oracle Data Provider, ODBC and OleDb via an Agent with appropriate drivers installed. If using 32-bit drivers, please use a 32-bit Agent.

In addition to 'SELECT', 'INSERT', and 'UPDATE' (depending on whether it's a source or a destination) permissions, the database user will also require the database role 'SELECT\_CATALOG\_ROLE' – this enables Eightwire to get a list of all relevant tables and columns and is required. Without this, you may simply not see one or more tables listed in Conductor that you were expecting to see, even if you have the other permissions required.

#### PostgreSQL

Supports ODBC and OleDb via an Agent with appropriate drivers installed.

#### Redshift

Supports direct cloud-to-cloud ODBC connections only. Utilises multiple simultaneous connections to achieve very fast data transfer speeds. Requires the *Database*, *Username*, *Password*, *Server*, *AWSRegion*, *Bucket*, *AccessKey* and *SecretKey* connection properties.

#### SQL Server

Supports SQL Native Client, ODBC and OleDb, either direct (e.g. SQL Azure) or via an Agent with appropriate drivers installed.

#### SYBASE

Supports ODBC and OleDb via an Agent with appropriate drivers installed.

### Binary Connectors

Most connectors extract structured data from the platforms they connect to. Binary connectors do not make any attempt to extract data from their platforms. Instead raw binary is read and written instead. This is represented with Conductor as structured tabular data with a binary component.



#### Folder








This connector allows file transfer within Conductor. This connector treats a Windows file system folder in the same way other connectors treat databases – as a container which holds tabular information.

The Data Store path points to a parent folder. This parent folder contains child folders. Each child folder contains files. This *parent-child-file* hierarchy is similar to a relational databases *database-table-row* hierarchy. When you browse the Data Store, you will see the list of child folders represented as tables. Each table will have the same columns – **Filename**, **Size**, **DateCreated**, **DateModified** and **Data**. You should recognise these as the columns you would see in Windows File Explorer when looking at one of those child folders. These columns are represented here as columns in a table. The *Data* column has a *binary* type and contains the binary content of each file in that child folder. Each row of data in the table comes from each file in the folder.

It is important to note that Conductor operates on a folder containing other folders, not simply on a folder containing files. This enforced two-level hierarchy of folders is to provide the same *database-table-row* structure that Conductor is designed to work with.

Consider the following folder structure and files as an example:

	Filename	Size	Created	Modified
	<b>Customers</b>			
	<b>Bobs Catering</b>			
	INV33201.pdf	9 KB	2018-07-13	2018-07-13

	INV34423.pdf	8 KB	2018-07-16	2018-07-17
	INV34555.pdf	10 KB	2018-07-17	2018-07-17
	<b>Kates Hardware</b>			
	INV19880.pdf	7 KB	2018-07-14	2018-07-14
	<b>Sams Software</b>			
	INV36310.pdf	13 KB	2018-07-12	2018-07-12
	INV36311.pdf	9 KB	2018-07-14	2018-07-17

To use this as a source Data Store, you would set the Data Store's *Path* to point to the *Customers* folder.

After Conductor has created the Data Store and scanned the folder, you will see the following *tables* in the Data Store:

```
Bobs Catering
Kates Hardware
Sams Software
```

Each of these tables has the same columns and data types as mentioned above.

To transfer all of the files in one of the child folders (e.g. *Bobs Catering*), you will need to build a Process for this child folder, which transfers files to a different destination *Folder* Data Store. When this Process executes it will read each file as a single row of data and write each row of data as a file into the destination *Folder* Data Store. For the *Bobs Catering* Process, the table of data that is being transferred would look like this:

<i>Filename</i>	<i>Size</i>	<i>DateCreated</i>	<i>DateModified</i>	<i>Data</i>
INV33201.pdf	9 KB	2018-07-13	2018-07-13	0xAE4F930017...
INV34423.pdf	8 KB	2018-07-16	2018-07-17	0xAE4F09F011...
INV34555.pdf	10 KB	2018-07-17	2018-07-17	0xAE4F5DE4FF...

It is of course possible to use a different type of destination Data Store and simply write this data into a database for some other purpose; or read from some other type of Data Store and write to file, provided all columns are mapped correctly, are non-null and of the correct types.

To filter the files that are copied, add one or filter conditions to the *Bobs Catering* table in the Data Store. For example, 'Filename is equal to "INV34423.pdf"'.

### Malware

*Conductor scans every file transferred for malware. Infected files will be rejected and not transferred. Conductor will scan inside ZIP and RAR files.*

***It is your responsibility to ensure that your files are not infected with viruses or other malware.***

This type of connector is not automatically made available to all accounts. If you cannot see it and wish to use it, please contact Eight Wire support to activate it.

## Appendix B - System Metadata

*System Metadata* is the metadata surrounding a Conductor Batch and the Process it's based on. A Batch is the one-time execution of a Process. The Batch does the work, while the Process defines what should be done. This system metadata may be added as one or more additional columns in a Process or referenced in any Expression.

This table gives the name of each specific metric, its meaning and what data type it returns:

<b><i>AllSystemData</i></b>	<p>A collection of all of the system data values listed here, in XML format.</p> <p><i>Returns XML.</i></p>
<b><i>BatchExecuteDate</i></b>	<p>The date and time this data was moved through Conductor. The time zone used is that of the Account that owns the Process.</p> <p><i>Returns a datetime in the format</i> <i>yyyy-MM-dd HH:mm:ss.ffffff</i></p>
<b><i>BatchExecuteDateUTC</i></b>	<p>The date and time this data was moved through Conductor in Universal Coordinated Time (UTC).</p> <p><i>Returns a datetime in the format</i> <i>yyyy-MM-dd HH:mm:ss.ffffff</i></p>
<b><i>BatchID</i></b>	<p>The internal Conductor ID for the Batch that moved the data.</p> <p><i>Returns a positive integer</i></p>
<b><i>BatchInitiatedBy</i></b>	<p>The name of the Conductor user or trigger that initiated data movement.</p> <p><i>Returns a text, such as:</i> <i>"JaneSmith" or "Scheduler"</i></p>
<b><i>DestinationAccountID</i></b>	<p>The internal Conductor ID for the destination Account.</p> <p><i>Returns a positive integer</i></p>
<b><i>DestinationAccountName</i></b>	<p>The name of the destination Account.</p> <p><i>Returns text</i></p>
<b><i>DestinationDataStoreID</i></b>	<p>The internal Conductor ID for the destination Data Store.</p> <p><i>Returns a positive integer</i></p>

<b><i>DestinationDataStoreName</i></b>	<p>The name of the destination Data Store.</p> <p><i>Returns text</i></p>
<b><i>ProcessID</i></b>	<p>The internal Conductor ID for the Process.</p> <p><i>Returns a positive integer</i></p>
<b><i>ProcessName</i></b>	<p>The name of the Process.</p> <p><i>Returns text</i></p>
<b><i>SourceAccountID</i></b>	<p>The internal Conductor ID for the source Account.</p> <p><i>Returns a positive integer</i></p>
<b><i>SourceAccountName</i></b>	<p>The name of the source Account.</p> <p><i>Returns text</i></p>
<b><i>SourceDataStoreID</i></b>	<p>The internal Conductor ID for the source Data Store.</p> <p><i>Returns a positive integer</i></p>
<b><i>SourceDataStoreName</i></b>	<p>The name of the source Data Store.</p> <p><i>Returns text</i></p>
<b><i>Version</i></b>	<p>The current Conductor version number.</p> <p><i>Returns text in the format x.x.x.x</i></p> <p><i>Generally the first part is the edition number and the rest indicates the date and build number. Every time Conductor is updated with new functionality the version number changes. For example:</i></p> <p><i>2.2017.111.1</i></p>

## Glossary of Terms

**Batch**

Used interchangeably with *job*, *data transfer* and *Process execution*. When you execute a Process, one or more Batches are created, each controlling and recording the progress of the data transfer defined by the Process.

**Conductor**

The proper name of the Eightwire data platform. It is also referred to as *data exchange*, *orchestrator*, and simply as *Eightwire*.

**Data Exchange, DX**

Usually refers to a specific industry sector use-case, community or network of Eightwire Conductor customers, united under a common purpose, such as a government sector. Conductor is frequently referred to as *The Data Exchange* in some security and risk documentation and in New Zealand government-related documentation.

**Orchestrator, Master Services**

Collectively, the Master Services cloud components of the Eightwire platform that control and 'orchestrate' the movement of data, but do not themselves handle data.

**Portal, Web Portal**

The subject of this User Guide. The web application used to manage Conductor Projects, Data Stores, Processes, etc. The web portal is part of Conductor's Master Services. Users do not need to remain logged in when running data transfers – once started data transfers (Batches) will continue even if the user logs out or closes their web browser.

**Processing Server, Processing Service**

One server node in a regional server cluster, designed to handle customer data under the command of the Master Services.

**Region, Processing Region**

Usually refers to a regional server cluster. Data transfers can optionally be configured to use one specific region when transferring data to ensure data sovereignty requirements are met.



## Contact Us

### Corporate

Email:	<a href="mailto:heya@eight-wire.com">heya@eight-wire.com</a>
Website:	<a href="https://www.eight-wire.com">https://www.eight-wire.com</a>
LinkedIn:	<a href="https://www.linkedin.com/company/eight-wire-limited/">https://www.linkedin.com/company/eight-wire-limited/</a>
Twitter:	<a href="https://twitter.com/eightwiredata">https://twitter.com/eightwiredata</a>
Phone:	(+64) 4 979 8838

### Support

Email:	<a href="mailto:support@eight-wire.com">support@eight-wire.com</a>
Intercom:	<i>Via the speech bubble in the lower right corner of the web portal</i>
FAQ:	<a href="https://www.eight-wire.com/help/">https://www.eight-wire.com/help/</a>