

iDashboards Data Hub Manual

Version 11.3

iDashboards Data Hub Manual

Version 11.3

No part of the computer software or this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from iDashboards. The information in this document is subject to change without notice. If you find any problems with this documentation, please report them in writing to support@iDashboards.com. iDashboards does not warrant that this document is error free.

Copyright © 2004 - 2022 iDashboards. All rights reserved.

Trademarks:

The iDashboards logo and tagline are trademarks of iDashboards.

All other products and company names referenced herein are the trademarks of their respective owners.

Support information:

iDashboards
900 Tower Drive, 4th Floor
Troy, MI 48098

Phone: (248) 528-7160

Fax: (248) 828-2770

Email: support@iDashboards.com

Web site: <http://www.iDashboards.com>


1. Table of Contents

1. TABLE OF CONTENTS	3
2. OSKAR	6
3. DATA HUB OVERVIEW	7
4. DATA HUB INSTALLATION	8
4.1 SYSTEM REQUIREMENTS.....	8
4.2 WINDOWS INSTALLATION	8
4.3 MANUAL INSTALLATION	8
4.3.1 <i>Deploying idbdata.war</i>	9
4.3.2 <i>Choosing a Context Root</i>	9
4.4 DATA HUB DATABASES	9
5. STARTING THE IDASHBOARDS DATA HUB APPLICATION	11
5.1 DETERMINING THE URL OF THE APPLICATION	11
5.2 LOGGING INTO THE IDASHBOARDS DATA HUB APPLICATION	12
5.3 THE DATA HUB HOME SCREEN	13
5.4 LOGGING OUT OF IDASHBOARDS DATA HUB APPLICATION	13
6. MANAGING USERS	14
6.1 THE IDASHBOARDS SYSTEM USER	14
6.2 ADDING A USER.....	14
6.3 DELETE OR MODIFY A USER	15
6.4 CHANGE PASSWORD.....	16
7. MANAGING GROUPS	17
8. EXTRACT, TRANSFORM AND LOAD	18
8.1 MANAGING DATA SOURCES FOR ETL.....	19
8.1.1 <i>JDBC and ODBC Drivers</i>	20
8.1.2 <i>Add a Data Source</i>	20
8.1.3 <i>Delete or Modify a Data Source</i>	21
8.1.4 <i>Use as Data Store Property</i>	21
8.1.5 <i>Using Stored Procedures</i>	22
8.2 DATA SETS.....	22
8.2.1 <i>Create Data Set</i>	22
8.2.2 <i>Table/View</i>	22
8.2.3 <i>Text File</i>	24
8.3 ETL JOBS.....	25
8.3.1 <i>Data Sets</i>	25
8.3.2 <i>Target Tables</i>	26
8.3.3 <i>Create New Target Table</i>	26
8.3.4 <i>Edit Target Table</i>	27

8.3.5	<i>Clone Table</i>	27
8.3.6	<i>Drop Table</i>	27
8.3.7	<i>ETL Job Scheduling</i>	28
8.4	ETL JOB STATUS	29
9.	IMPORTED DATA SOURCES	30
9.1	USAGE REQUIREMENTS	31
9.2	CONFIGURE THE IMPORTED DATA DATABASE	31
9.3	CREATING NAMED RANGES IN AN EXCEL WORKSHEET	33
9.3.1	<i>Header Rows in Named Ranges</i>	34
9.4	IMPORTED FILES	35
9.4.1	<i>File Name Rules</i>	37
9.5	AUTOMATED FILE UPLOADS	38
9.5.1	<i>Enable the setting</i>	38
9.5.2	<i>Auto Uploader System and User Requirements</i>	38
10.	SYSTEM	39
10.1	MODIFYING A SYSTEM SETTING	39
10.2	SETTINGS	39
10.2.1	<i>Server Settings</i>	40
10.2.2	<i>Security Settings</i>	41
10.2.3	<i>SMTP Settings</i>	41
10.2.4	<i>Notification Email Settings</i>	42
10.3	AUTHENTICATION SETTINGS	42
10.4	API ACCOUNTS	42
10.5	CONNECTED APPS	42
10.6	IMPORTED DATA CONFIGURATION	43
10.7	IMPORTING AND EXPORTING DATA SETS AND ETL JOBS	43
10.7.1	<i>Criteria for Export</i>	44
10.7.2	<i>Criteria for Import</i>	45
10.7.3	<i>Managing Global Identifiers (GIDs)</i>	45
10.7.4	<i>Exporting Data Sets and ETL Jobs</i>	47
10.7.5	<i>Importing Data Sets and ETL Jobs</i>	49
10.8	LOG SETTINGS	52
10.8.1	<i>Downloading Log Files</i>	53
10.8.2	<i>Sending Log Files to iDashboards Technical Support</i>	53
10.8.3	<i>Log Viewer</i>	54
10.8.4	<i>Log Configuration</i>	54
10.9	SERVER STATUS	55
10.9.1	<i>Pausing and Restarting the Server</i>	55
10.9.2	<i>Understanding Server Events</i>	56
10.9.3	<i>Troubleshooting Error Events</i>	57
10.9.4	<i>Event Retention</i>	57
10.9.5	<i>Email Events</i>	58
10.10	VIEW MANUAL	58
10.11	CHANGE PASSWORD	58

10.12	UPLOAD LICENSE	59
10.13	CHANGING LANGUAGE	59
10.14	ABOUT	59

2. OSKAR



OSKAR, the Online Support & Knowledge Acquisition Repository, is the preferred support resource for iDashboards' customers, partners and prospects. The OSKAR Support Portal can be used to submit, review and update support tickets.

<https://oskar.idashboards.com/>

Those who have an active, support and maintenance contract with iDashboards also have access to the following content in our User Community:

- **Knowledge Base** – Numerous product and technology articles for your review.
- **Community** – Forums and discussion groups for customers to discuss various topics and products amongst themselves.
- **Resources** – Many downloadable resources that can be used with iDashboards.
- **Ideas** – Area for customers to submit feature requests and great product ideas.
- **Blog** – Thoughts, stories and ideas on data and dashboards

3. Data Hub Overview

iDashboards must be installed and licensed for the Data Hub to function. The 2 main components of the iDashboards Data Hub are:

- 1) Data Connection/Manipulation
 - a) Data Sources
 - b) Data Sets
 - c) Identifying Data Stores and Target Tables
 - d) Transforming data with ETL Jobs (Extract, Transform, and Load)
 - e) Scheduling ETL Jobs
- 2) Imported Data Database
 - a) Configure, Connect and Initialize the database
 - b) Upload Excel, Delimited and Fixed Column files
 - c) Use within the iDashboards User Application

4. Data Hub Installation

The iDashboards Data Hub is an application focusing on the hosting and manipulating of data. This application is intended to assist with complex data-related tasks which are not always possible within the iDashboards User Application. There are two main components of the iDashboards Data Hub: Data Manipulation via ETL and a special Imported Data Database for optimizing the way companies' use Excel, Delimited and Fixed Column files as a data source within iDashboards.

The iDashboards Data Hub is intended to be installed on, or within the same network as, the installation of the iDashboards application server. Please refer to the iDashboards Administrative Manual for details.

4.1 System Requirements

- **iDashboards Application Server** – An existing iDashboards v9.5 or later installation
 - **System Requirements** – Refer to the iDashboards Administrator's Manual
 - **License** – The Data Hub requires licensing which may require a new license file if upgrading. Contact your Account Manager to verify if you have a valid license for the Data Hub.
- **Databases and user privileges**
 - Supported databases include: SQL Server (2005 or later), Oracle (9i or later), MySQL (5.0.3 or later) and PostgreSQL (9 or later). These types of databases are required for the ETL functionality and the Imported Files database functionality.
 - A database user with read privileges to the ETL data store database and Imported Data database
 - A database user with write, create table, update table, and drop table privileges to the ETL data store database and Imported Data database.
 - You can combine all privileges into one user, or use one user for read access and another user for write, create table, and drop table privileges.

4.2 Windows Installation

New installations of iDashboards include the Data Hub. Refer to the Administrators Manual for additional details.

4.3 Manual Installation

The Data Hub is included with new manual installation packages. Refer to the Administrators Manual for additional details on installing iDashboards.

Note: The installation process assumes that iDashboards has already been fully installed.

Note: All installation files are located in the bin directory on the iDashboards installation media.

The exact procedure for installing iDashboards Data Hub may vary from one installation to the next. The process normally involves copying a WAR file into a specific directory, copying necessary files, and then restarting the server.

4.3.1 Deploying `idbdata.war`

Locate the file “`idbdata.war`” and copy it into the application server deployment folder. Another option includes uploading the WAR file through the application server using a web interface.

Example: `...\Tomcat\webapps\idbdata.war`

4.3.2 Choosing a Context Root

Regardless of the application server used to host the Data Hub, it must be assigned a “context root” within the server’s URL space. Conceptually, the context root can be thought of as a subdirectory beneath the server’s root URL, which forms the root of the web application’s URL space. This allows multiple web applications from different sources to be deployed to the same application server without URL conflicts.

It is recommended that “`/idbdata`” be used as the context root for the Data Hub web application. Since the Data Hub WAR file is named `idbdata.war`, some application servers (such as Tomcat) will automatically default its context root to “`/idbdata`”, so choosing that can simplify the deployment process.

4.4 Data Hub Databases

In addition to the iDashboards Repository Database (refer to the iDashboards Administrator’s Manual for details), the Data Hub requires a minimum of two additional databases for the ETL Data Store and Imported Data.

- Supported databases include: SQL Server (2005 or later), Oracle (9i or later), MySQL (5.0.3 or later) and PostgreSQL (9 or later). If you have not yet installed the necessary JDBC driver for your desired database, refer to the iDashboards Administrator’s Manual for driver installation details.
 - Create two new databases
 - One database for Imported Data. This may be named anything. Best practices for naming the database is ***idb_imported_data***
 - A minimum of one database for an ETL Data Store. This may be named anything. Best practices for naming the database is ***idashboardsdatastore***
 - Prepare a database user with read privileges to the ETL data store database and Imported Data database, and a database user with write, create table, update table, and drop table privileges to the ETL data store database and Imported Data database.

- You can combine all privileges into one user, or use one user for read access and another user for write, create table, and drop table privileges.
- Once the databases are created and the Data Hub is installed, you may configure the databases for use from within the Data Hub application.
 - Refer to Chapter 8 “Extract, Transform and Load” to configure the ETL Data Store database
 - Refer to Chapter 9 “Imported Data Sources” to configure the Imported Data database

5. Starting the iDashboards Data Hub Application

The iDashboards Data Hub is a browser-based application used to configure data tables and a special Imported Data database. It is available once the `idbdata.war` file has been deployed to the application server and properly configured. Some key tasks that may be performed through the iDashboards Data Hub are:

- Creating, Running and Scheduling ETL Jobs
- Uploading Microsoft Excel, Delimited and Fixed Column files

5.1 Determining the URL of the application

The URL of the Data Hub is the web address used to access settings with a web browser. To determine it, you must know three things:

1. The hostname or IP address of the server on which iDashboards Data Hub is deployed.
2. The port number on which the iDashboards Data Hub application server is listening, if it is other than port 80, which is the default HTTP port number.
3. The “context root” under which the iDashboards Data Hub application was deployed. Normally this will be “idbdata.” See Section 4, “Data Hub Installation” for more information.

Once these three components are known, the URL of the Administrator Application will be:

```
http://<servername or IP address>:<port number>/<context root>/idbdata
```

For example, if the hostname is “dashmachine”, the port number is 8080, and the context root is “idbdata”, the URLs for iDashboards would be:

Interface	URL
iDashboards Application	<code>http://dashmachine:8080/idadashboards</code>
▶ Data Hub Application	<code>http://dashmachine:8080/idbdata</code>
<i>If the port number is 80, it can be omitted from any of the URLs, as seen below:</i>	
▶ Data Hub	<code>http://dashmachine/idbdata</code>

▶ = Topic included within this manual

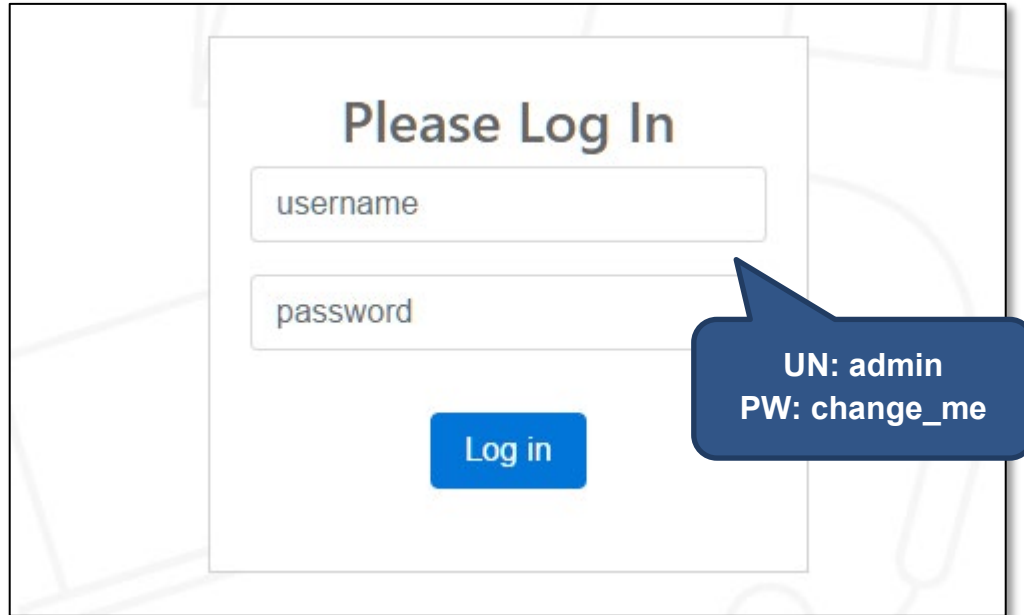
In some cases, such as when the server is being accessed over the Internet, it may be necessary to append a domain name to the hostname:

```
http://dashmachine.mycompany.com/idbdata
```

Once the iDashboards Data Hub URL is accessed through a web browser, the login screen should appear.

5.2 Logging into the iDashboards Data Hub Application

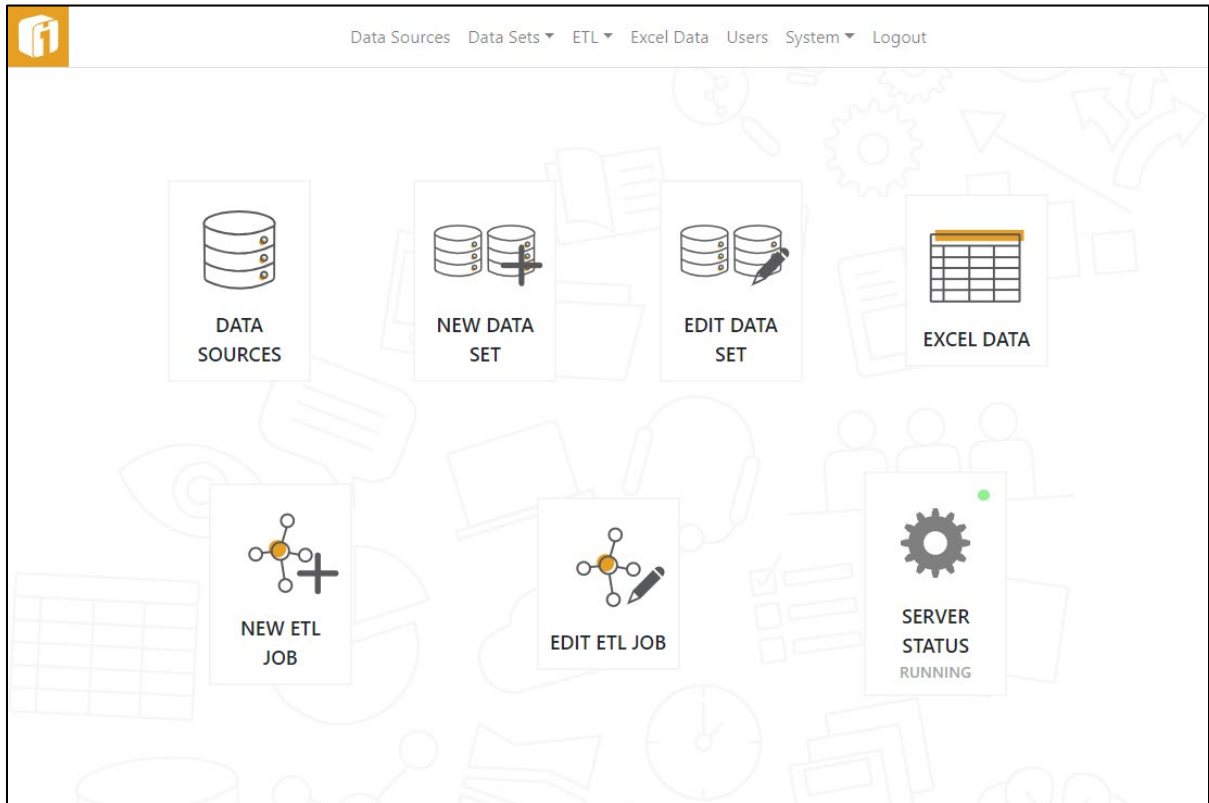
When iDashboards is initially installed, an administrative user account is created with the username *admin* and the password *change_me*. These credentials can be used to log into the Administrator Application through the login screen.



Note: It is highly recommended that the admin password is changed after installation.

5.3 The Data Hub Home Screen

Upon successful login, the Data Hub Home screen will be displayed. Icons on the Home screen are popular navigation options, whereas full navigation appears along the top of the page.



5.4 Logging out of iDashboards Data Hub Application

To log out of the iDashboards Data Hub Application, click "Logout" on the menu bar.

6. Managing Users

Users can be entirely managed via the iDashboards Administrator Application. Users can be partially managed via the iDashboards Data Hub Application with the exception of Groups and Data Source Access Control.

Each user requiring access to the iDashboards Data Hub must have a user account with the role of either “**Admin**” or “**Data Admin**”. Users can be viewed, edited, created and deleted through either application if the users’ role is “Admin”. However, if the users’ role is “Data Admin”, the “Users” page is not available.

Please refer to the iDashboards Administrative Manual for details.

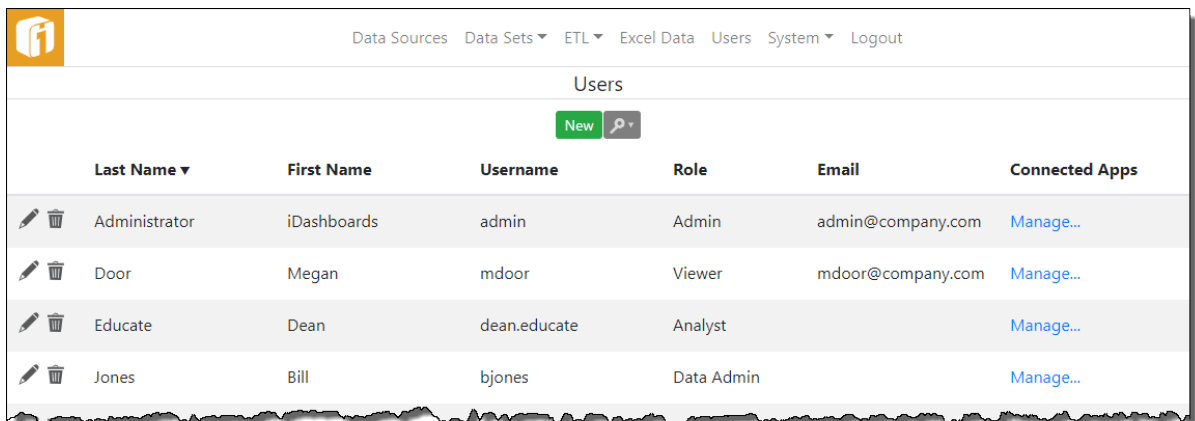
6.1 The iDashboards System User









When iDashboards is installed, a user account with the username “admin”, the password “change_me” and the Admin role is created. This account, which cannot be deleted, is referred to as the “iDashboards system user.” All of the properties for this user account, except the username and role, can be changed. **You must be logged in as the iDashboards system user to edit any information about the iDashboards system user.**

Note: It is highly recommended that the “admin” password is changed after installation.

6.2 Adding a User

New user accounts are created in the “Users” screen of the Data Hub Application. Click “Users” on the application menu to view up to 1000 users. The default sort order of the list is by “Last Name”. Click any column header to sort by that column.



	Last Name ▼	First Name	Username	Role	Email	Connected Apps
 	Administrator	iDashboards	admin	Admin	admin@company.com	Manage...
 	Door	Megan	mdoor	Viewer	mdoor@company.com	Manage...
 	Educate	Dean	dean.educate	Analyst		Manage...
 	Jones	Bill	bjones	Data Admin		Manage...

To create a new user, click the “New” button and then enter the user’s first name, last name, user name (login ID) and password into the appropriate textboxes. Select a role for the user from the “Role” dropdown.

Note: Assigning and identifying “Group” permissions is only permissible in the iDashboards Administrator Application and not the iDashboards Data Hub.



- **Viewer** – Users assigned to this role have read-only access to dashboards within the categories where they have permissions. They can, however, fully interact with charts and dashboards. They do not have a “Personal” category and there are limited menu options.
- **Builder** – Users assigned to this role have full dashboard, chart and picklist building capabilities within the categories where they have “save” permissions. They can upload files to the “Content” directory and they do have a “Personal” category.
- **Data Admin** – This role allows all Builder role functions, plus the ability to log into the Data Hub for performing data manipulations and Excel, Delimited and Fixed Column file uploads.
- **Admin** – Users assigned this role have all the permissions of the Data Admin role and are the only users that can log into the Administrator Application and perform administrative functions.

Once the basic user information has been entered, select the “Save” button to add the new user. The newly-added user will appear in the list, with location based on the sort-order in effect.

Values entered for a new user must adhere to the following rules:

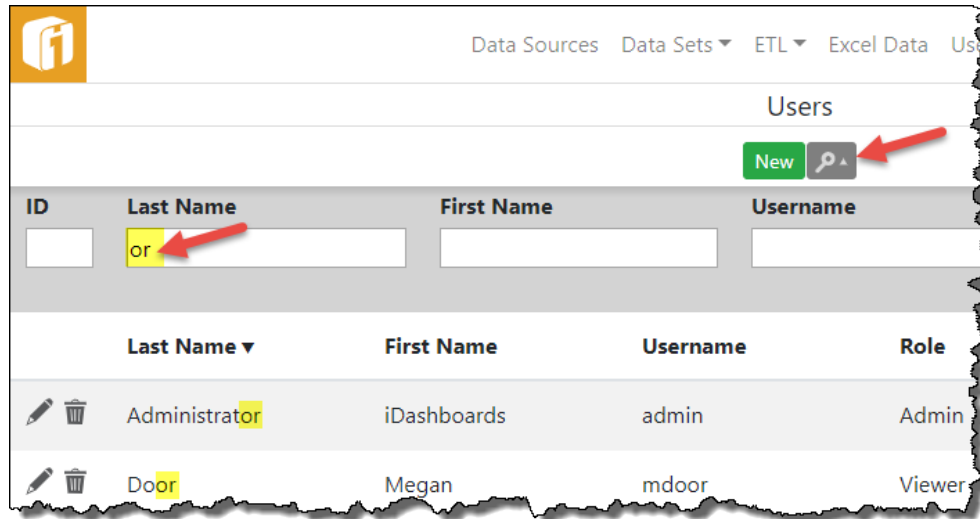
- **Last name** — 1 to 50 characters.
- **First name** — 1 to 50 characters.
- **Username** — 1 to 75 characters. The username must be unique within the system, without regard to case. In other words, “ADMIN” would be considered the same as “admin.” However, when logging into the system, a user should assume the username is case-sensitive.
- **Password** – 1 to 50 characters. Non-whitespace characters. Passwords are always case-sensitive.
- **Email** — (optional) Up to 100 characters.

6.3 Delete or Modify a User

Any user account except the iDashboards system user can be deleted. To delete a user account, click its Delete icon (). To modify a user account, click its Edit icon (). The edit user screen will appear through which the user’s basic information can be edited.

To change a users’ password, first edit the user account and then check the box for “Change Password”. Note that a user’s password will NOT be changed unless values are entered into both the “Password” and the “Confirm Password” textboxes, the values match, and they are a legal password (4-15 non-whitespace characters.)

To locate a user account, click the search icon to show the various fields in which to search against. Note that in certain cases of troubleshooting, searching for a user ID can be performed here.



6.4 Change Password

Users with the role of "Admin" or "Data Admin" can change the password of the currently-logged-in user via:

Click "System" > "Change Password"

7. Managing Groups


Group permissions are not available within the Data Hub administration. Defining Data Access control, via groups, is not possible within the Data Hub administration. Data Source Access Control is not available on Excel, Delimited and Fixed Column files uploaded via Data Hub - neither at the file level nor Imported Data database level.

Please refer to the iDashboards Administrative Manual for details.

8. Extract, Transform and Load

This function of the Data Hub is purposed to alter the original Data Set into a modified data set using a series of incremental changes along the “data flow”. The original data can go through numerous transformations, ranging from complex to simple, until the desired output data set is achieved. A series of Extract (E), Transform (T) and Load (L) tasks will perform the work using a graphical interface and flow lines between each task. Each “data manipulation project” is referred to as an ETL Job.

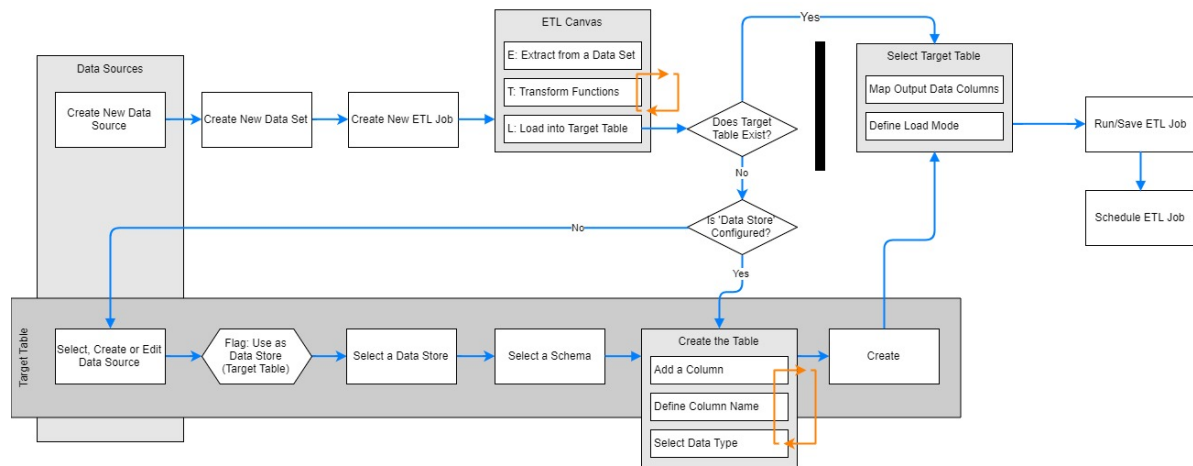
Before designing a new ETL job, it is necessary to configure certain areas of the Data Hub. For discussion purposes, this illustrates one popular method for planning the first ETL job within the Data Hub; with an outline of the same process.



1. Create a new Data Source
 - a. Use an JDBC driver
2. Create a new Data Set
 - a. Connect to a Data Source and define the query
3. Begin a new ETL Job
 - a. Extract from a Data Set
 - b. Transform the data
4. Create a new Target Table
 - a. Create/Select/Edit a Data Source
 - b. Enable as a Data Store
5. Complete the ETL Job
 - a. Load data output into a Target Table

The remainder of this chapter will discuss each of these steps.

First Time ETL Diagram



8.1 Managing Data Sources for ETL

Refer to the **Administrators Manual** for details on **Managing Data Sources**. **Data Sources** can be equally managed between the **iDashboards Administrator Application** and the **iDashboards Data Hub**. However, the context of this manual is to describe **Data Sources** for the purpose of **ETL**.

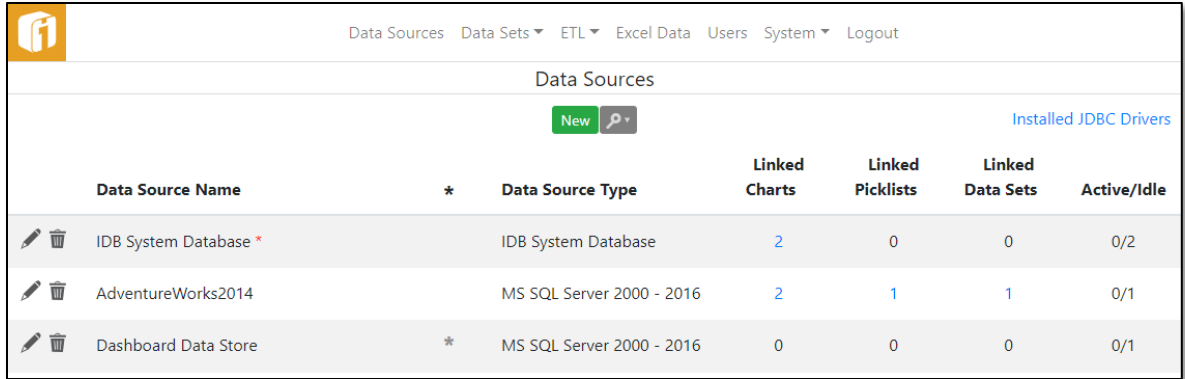
Migrating an existing instance of iDashboards to use the Data Hub will simply replicate the same Data Sources between both applications. Adding, Editing or Deleting Data Sources can occur using either application.







For the purpose of ETL, Data Sources are necessary for two reasons:

- 1) Establishes a source for creating Data Sets (required for the Extract task)
- 2) Establishes a connection into a production database for the purpose of
- 3) Required for identifying and configuring Target Tables (used during the Load) by way of identifying a Data Source as a “Data Store”

Configuring a Data Source usually requires special knowledge about the target database, for example usernames, passwords, port numbers, host names, drivers, etc., so it generally should be done by, or with the assistance of, a database administrator (DBA).

Data Sources are managed through the Data Sources screen, which is accessed by clicking “Data Sources” in the application menu. The Data Sources screen displays a list of currently configured data sources. If one of the listed Data Sources is the iDashboards repository database, it will be marked with a red asterisk (*) after its name. The default name it is given during the iDashboards installation process is “IDB System Database”, although this can be changed.



Data Sources							
Data Source Name		*	Data Source Type	Linked Charts	Linked Picklists	Linked Data Sets	Active/Idle
 	IDB System Database *		IDB System Database	2	0	0	0/2
 	AdventureWorks2014		MS SQL Server 2000 - 2016	2	1	1	0/1
 	Dashboard Data Store	*	MS SQL Server 2000 - 2016	0	0	0	0/1

8.1.1 JDBC and ODBC Drivers

The Data Hub does not allow for ODBC drivers/connections. This does not restrict the use of Microsoft Excel, Delimited and Fixed Column files when uploading into the configured Imported Data database.

To view the installed JDBC drivers, click the link “Installed JDBC Drivers” from the “Data Sources” page.



Please refer to the iDashboards Administrative Manual for details.

8.1.2 Add a Data Source

Once the required JDBC drivers have been installed, a Data Source may be configured. To add a Data Source select the “New” button to launch the “Create a Data Source” screen.

Selecting different Data Source Types may result in varying fields compared to the image below. If the field is required by iDashboards, it will be identified with an asterisk. However, specific data sources may have additional field requirements not designated with an asterisk.

8.1.3 Delete or Modify a Data Source

Any data source except the iDashboards repository can be deleted. To delete a data source, click its Delete icon (). To modify a data source, click its Edit icon (). The data source screen will appear through which the user's basic information can be edited.

8.1.4 Use as Data Store Property

For an ETL job to “write” into a database table, the Data Source needs to be configured as a Data Store. Enabling this option will require a “Write User” and “Write User Password” for securely controlling this permission. Only the supported databases listed below can be flagged as a Data Store:

- Microsoft SQL Server (2005 or later)
- Oracle (9i or later)
- MySQL (5.0.3 or later)
- PostgreSQL (9 or later)

8.1.5 Using Stored Procedures

Refer to the Administrators Manual for details on Stored Procedures. Stored Procedures can be equally managed between the iDashboards Administrator Application and the iDashboards Data Hub, with variations only in appearances.

To begin the process of configuring Stored Procedures from within the Data Hub, edit a configured data source to access the “Stored Procedure” option.

Note: Stored Procedures are only available with certain types of databases. Currently, stored procedures may only be used with Microsoft SQL Server, MySQL and Oracle databases.

8.2 Data Sets

Data Sets are required for the Extract “E” process of ETL. They include all of the definitions for querying a data source in preparation of transforming the data.

8.2.1 Create Data Set

From the home screen, click “New Data Set”; alternatively, from the navigation, Click “Data Sets” > “New”. Then, from the left-hand menu, select a data set option. For illustration purpose, only the “Table/View” option is outlined in the following section; however, each option is similar in concept.

Note: For information about creating a Data Set during an ETL Job configuration see ETL Jobs’ Section 8.3.1, “Data Sets”.

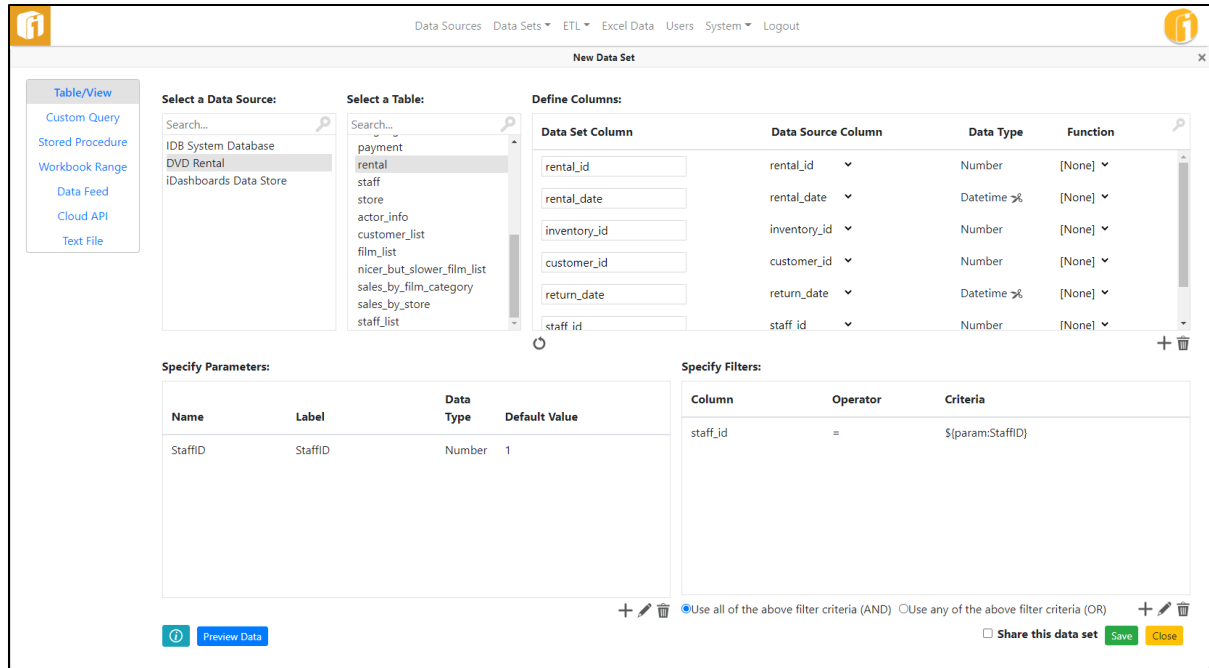
Note: The “Cloud API” option is documented within the product and can be found beneath “System” > “API Accounts”.

8.2.2 Table/View

This option allow a data set to dynamically connect to existing Tables or Views. The minimum requirements are to select a Data Source and then a Table or View.

Select a Data Source – Click on an item in the Data Source list, and the Table or View options will appear.

Select a Table or View – Here, a list of database schemas within that data source may appear. If schemas appear, click to expand/collapse to display the table or view names. Lastly, click on an item in this list, and the remaining Data Set fields will appear.



Define Columns – Within this section, data set columns can be:

- Renamed
- Reordered (via drag-n-drop)
- Remapped (via “Data Source Column” drop down menu)
- Applied a function (via “Function” drop down menu)
- Manually added (via “+” icon)
- Manually deleted (via “🗑️” icon)

Specify Parameters – Input Parameters provide the ability to define a variable when querying the data source. The input values are used during the data set definition or from the ETL jobs in a variety of ways, including:

- Filtering the data set based on the input values
- Using the input values in a chart’s custom SQL queries
- Using the input values in a column expression
- Get defined by each ETL job that extracts from the data set

Parameters can be:

- Added (via “+” icon)
- Edited (via “✎” icon)
- Deleted (via “🗑️” icon)

Name – This is the name to be used internally to identify the parameter. These values can be used within filter or expression dialogs. Parameter names must begin with a letter, consist of letters, numbers and underscores, and be no longer than 25 characters. Blank spaces are not allowed.

Label – This is the label displayed across the interface when the technical name is not required. Therefore, the label should be user-friendly and formatted nicely. It can contain up to 50 printable characters without leading or trailing whitespaces.

Data Type – This can be one of the three standard iDashboards data types; String, Number or Datetime.

Required Flag – This indicates whether or not a value is required for the Input Parameter before the chart or dashboard is displayed. If it is required, and a value is not present via a drilldown or initial value, the user is prompted to enter a value.

Replace blank with (optional, charts only) – If a user submits an Input Parameter with a blank value (i.e. they don't enter a value but still submit), the value in the "Replace blank with" field will be used instead.

Note: "Blank Parameter Values" can cause problems.

Initial Value (optional) – This is the value that will be used for the Input Parameter if no other value is present.

Specify Filters –

Filters can be:

- Added (via "+" icon)
- Edited (via "✎" icon)
- Deleted (via "🗑" icon)

Data Set Description – Click the "I" button (in the lower left) to rename the Data Set or add a description regarding the Data Set (500 MAX characters). This description can be viewed by clicking the "I" button again or from within the menu for editing a Data Set.

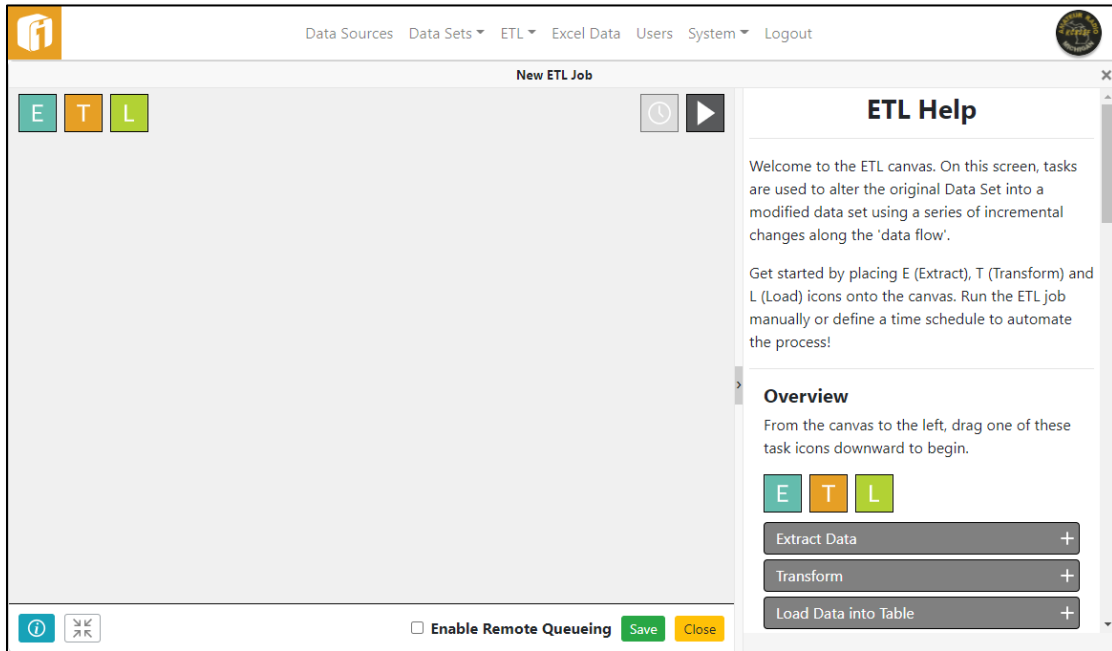
Note: Refer to the User Manual chapter "Advanced Chart features" for additional details.

8.2.3 Text File

This type of data set is created by referencing one, or more, text files. A common use-case would be using alike log files from anywhere on the Data Hub's server. The default location for text files in this application's '..\vizgroup\text_data\' folder. Other folders on the server can be added by creating the EXPOSED_FOLDERS.TXT file in '..\vizgroup\config\', and populate it with the folder's full path. When creating the data set the default and exposed folders are available for selection. Only one folder can be used, but multiple files in that folder can be selected.

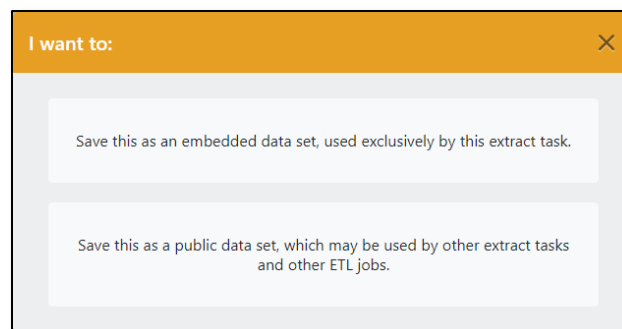
8.3 ETL Jobs

Click on “New ETL Job” from the Home screen or from the “ETL” navigation. Here, the ETL canvas will be displayed. When nothing is selected with the mouse-cursor, the ETL Help dialog will appear and offers in-product documentation on each of the tasks (right-hand side of the screen). Once the ETL Jobs are populated, this area of the screen will display a “One-Click” summary of the selected task.



8.3.1 Data Sets

A Data Set, can be create during the configuration of the Extract task. A Data Set create this way can also be saved as either a public Data Set, or an Embedded Data Set. A public data set may be used by other extract tasks and other ETL jobs. An embedded data set is save for exclusive use by its extract task, and will be named “<embedded:x>”, where “x” is the task number associated with its Extract Data task. The purpose for embedded data sets is to avoid clogging up “public” data sets when an ETL job has a one-time need for a data set. Extracts be changed between public Data Sets and Embedded Data Sets by using “Save As” when its configuration is completed.



8.3.2 Target Tables

Target tables are required for the end of an ETL diagram. It is the output location for ETL jobs where data is written into a database table; which is referred to as a “Data Store”. Refer to Section 8.1.4 “Use as Data Store Property” for information on setting up Data Stores. If the configured Data Source (flagged as a Data Store) does not already have a Target Table for the particular ETL job, then the target table can be created from within the Data Hub.

Understanding the “Target Table” features is equally comparable to manually creating, editing and dropping database tables through the database management utility. This feature, however, allows tables to be created, edited and dropped through the Data Hub interface for ease of use.

8.3.3 Create New Target Table

Creating a new Target Table is especially easy when the ETL job is fully configured because the Output Columns from the last task will help identify the database table that needs to be created! The Target Table can be auto-generated (preferred) or manually created.

8.3.3.1 Auto-Generated

- This method only requires the entry of a table name
- How it works: The last transformation function determines the details for the output data columns that will be automatically created in the desired Data Store. This includes the number of columns and column data type. During the last step of this process, the user can override the Column Name and Data Type, Column order, set the option to allow nulls and even create or remove columns in the table.

8.3.3.2 Manually Created

- Target Tables can be created from within the Data Hub by navigating to “ETL > New Target Table”. Alternatively, the database table can be created using database tools.
- How it works: The last transformation function determines the details for the output data columns that will need to be manually created - this includes the number of columns and column data type.

From any screen (including the ETL Job screen), select “ETL” from the navigation and then “New Target Table”. Follow the steps for creating the table by selecting the Data Store, the database schema (sometimes optional), and then create all of the data columns using the button “Add Column”.

Column Name – Enter a unique column name that will eventually map back into the ETL job

Data Type – Each database brand has unique data types, but the most universal options will be displayed when the dropdown list appears. To see additional data types click “More...”

Allow Nulls – This flag is required by database tables and therefore offers the control to allow or not-allow NULL values in the data column. By default, the value is set to Allow NULLs.

PK – Flags the field as part of the table’s primary key.

Index – Options are “None”, “Unique”, or “Non-Unique”.

Note: It is recommend, and required for Data Forms (if enabled), that a table has one unique primary key. It can be a single column, or multiple columns collectively identifying a unique record in the table.

8.3.4 Edit Target Table

CAUTION: Updating a table will destroy all data in that table.

From the navigation, select “ETL” and then “Edit Target Table”. Progress through the windows to locate the desired target table and then add/remove/edit the data columns until the new structure is desired. It is not possible to rename the target table.

Editing a Target Table can also occur when utilizing the Auto-Generated target table against a pre-existing table.

8.3.5 Clone Table

From the navigation, select “ETL” and then “Clone Table”. This task will make a new table based on the exact copy of an existing table. Elements that are cloned include:

- Column Name
- Data Type
- Allow Nulls setting

Note: Cloning a table does not copy any of the data. Therefore, after cloning a table, the new table will be empty.

8.3.6 Drop Table

CAUTION: Dropping a table will destroy the table. All data will be removed. All columns will be removed.

From the navigation, select “ETL” and then “Drop Table”. This task will completely delete (drop) a database table, including the removal of all records of data, associated columns and column data types.

8.3.7 ETL Job Scheduling

The schedule determines how often the ETL Job will be executed. Schedule frequencies range from every minute to once per year. More details can also be found in ETL Help's in-product documentation under "Schedule ETL Job".

Schedule this Job [X]

Next Run Time:

Months: No months were specified
Dates: No dates were specified
Times: No times were specified

January	February	March	April
May	June	July	August
September	October	November	December

Enabled? Save Cancel

Use CTRL+Click to select multiple items, individually or while holding down click draw a box over the items and release.

Schedule this Job [X]

Next Run Time:

Months: No months were specified
Dates: No dates were specified
Times: No times were specified

January	February	March	April
May	June	July	August
September	October	November	December

Enabled? Save Cancel

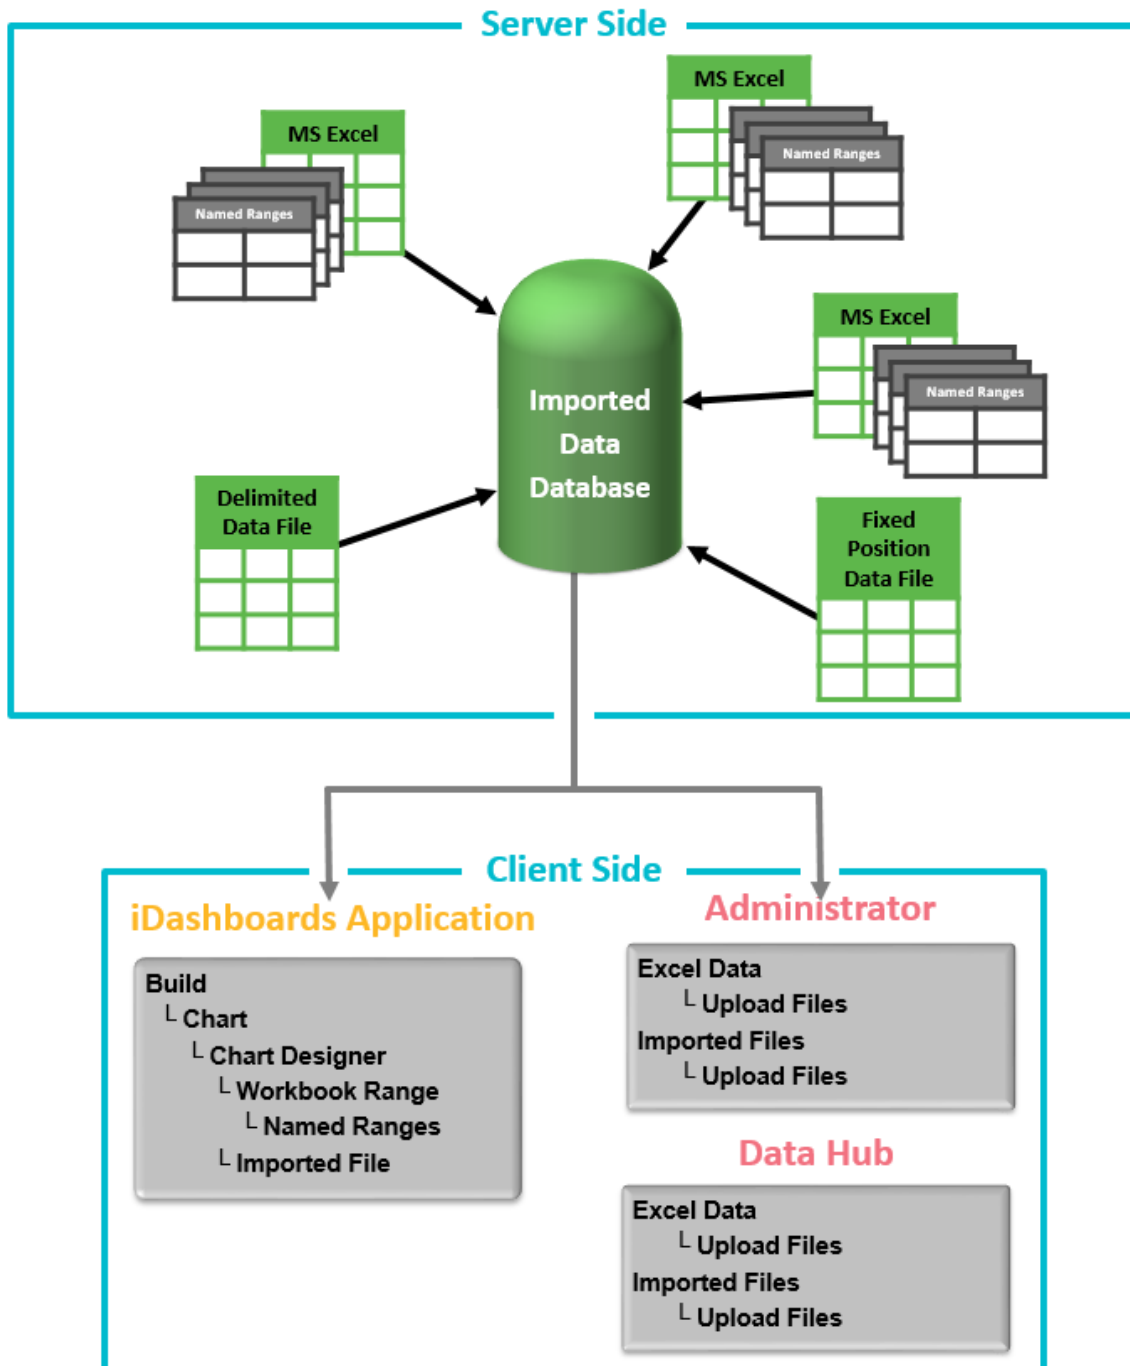
8.4 ETL Job Status

From the navigation, select “ETL” and then “Job Status”. Here, the history of executed ETL jobs will appear. The primary purpose of this screen is to provide troubleshooting on failed jobs or to diagnose why data tables might be changing (because the ETL job is scheduled to change the data). To help locate specific items, click the search icon and enter criteria to narrow down the list to a specific set of results.

To pause or resume all scheduled jobs, refer to Section 10.5, “Server Status”.

9. Imported Data Sources

Microsoft Excel, and Delimited and Fixed Position, data files can be equally managed between the iDashboards Administrator application and the iDashboards Data Hub. However, the context of this manual is to describe how to use these files for building charts within iDashboards. All files will be uploaded into the singular Imported Data Database.



9.1 Usage Requirements


- 1) Configured Imported Data database
- 2) **Excel** files
 - a) *.xls, *.xlsx files
 - b) Must contain named ranges
 - c) Macros (*.xlsm) are not supported
- 3) Data files
 - a) **Delimited**
 - i) *.csv files
 - ii) The Default is comma delimited
 - iii) Uses an optional column header row
 - b) **Fixed** Position
 - i) *.txt files
 - ii) Each column is in a specific position in the record
 - iii) Does not use a column header row, columns are defined during import
- 4) Duplicate Column Names are not allowed
- 5) File size must be less than 10MB

9.2 Configure the Imported Data Database

Before you can begin uploading Excel or Data Files to the Data Hub, you must first have a configured Imported Data database. The database must first be created on your database server and a user with proper permissions should be available. See “Data Hub Databases” in Chapter 4.4 for more information.

Once you have created your Imported Data database, follow the steps below to configure the database for use.

1. Log on to the iDashboards Data Hub application, from the System menu, select “Imported Data Configuration” a screen will display the configuration status.



2. Select “Edit Configuration” to open the Configure Imported Data Database. Enter the information required to connect to your newly created database. All fields marked

with an asterisk are required by iDashboards, but your database server may require additional information to make the connection.

3. Once you have entered your connection information and select “Save”.
4. If the database was configured correctly, you should see a green checkmark next to all of the status lines in the Imported Data Database Status screen.

You can now use “Imported Data” to choose “Excel Data” or “Imported File”, and then use the “Upload Files” button.

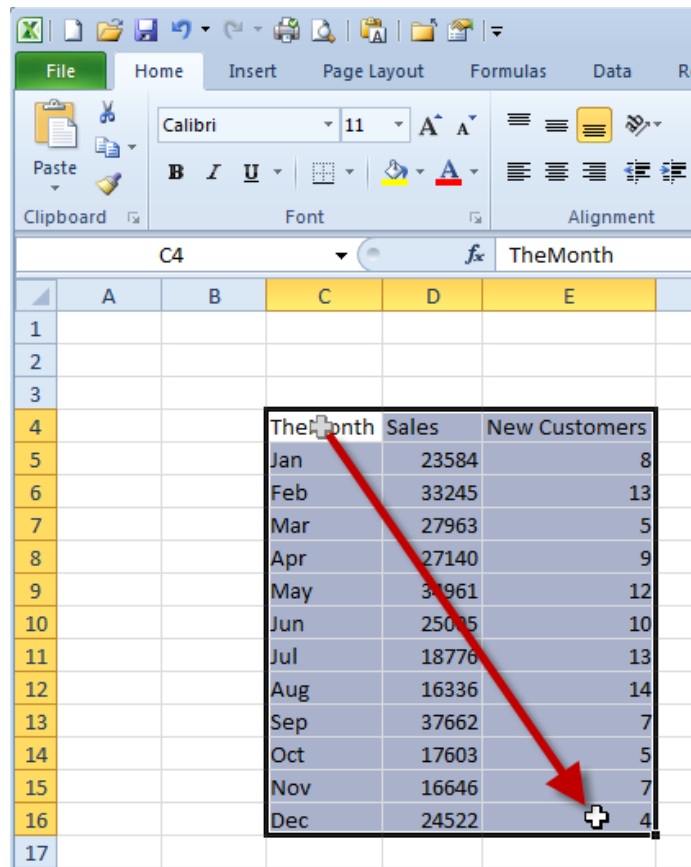
9.3 Creating Named Ranges in an Excel Worksheet

For iDashboards to use an Excel spreadsheet as a dynamic data source, the spreadsheet must contain one or more named ranges. A named range is an Excel technology that gives a name to a block of cells in the spreadsheet. A spreadsheet can have any number of named ranges and, as far as iDashboards is concerned, it doesn't matter what worksheet they are located on. When the Excel file is uploaded, each named range within the spreadsheet will be written into a database table. When you create a chart and select "Workbook Range" as the data source type, you will see the list of uploaded files and the associated named ranges. Select the Named Range that contains the data you would like to display in the chart and continue the chart building process.

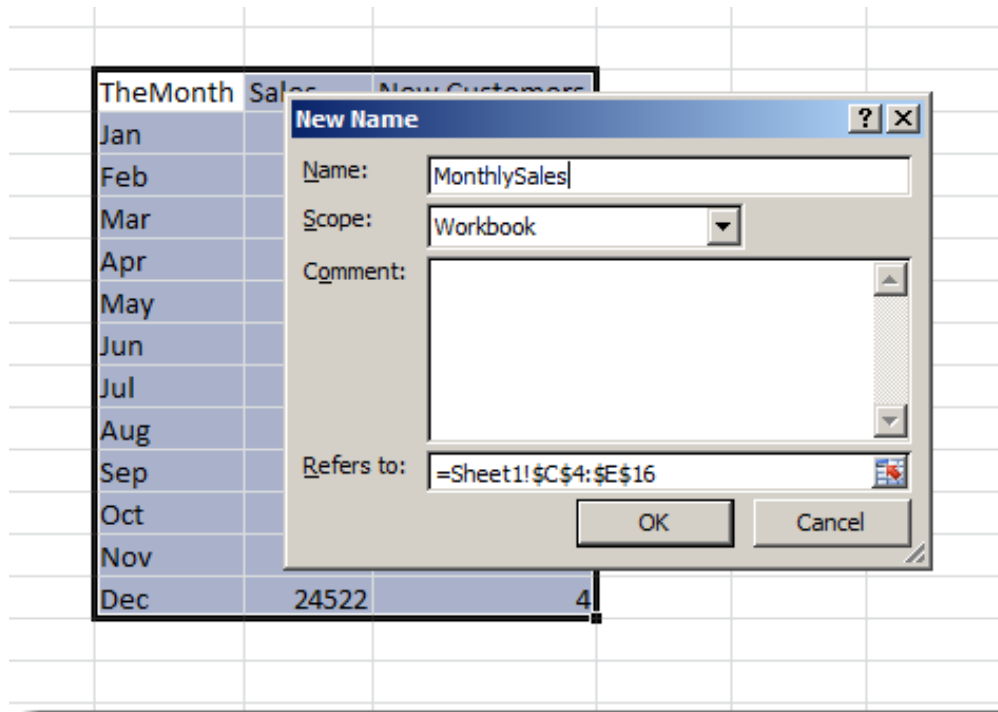
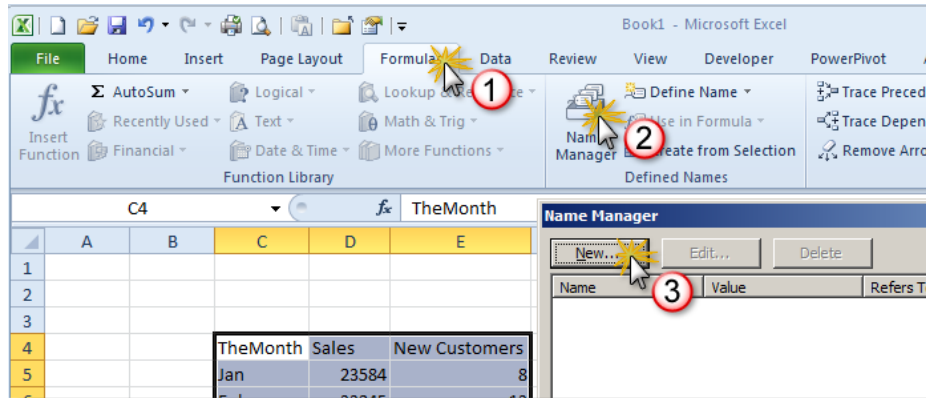
Tip: To learn more about named ranges, navigate to "Excel Data" and click on the button "Excel Named Range Tutorial".

The following steps are used to create a named range in an Excel workbook:

1. Use the mouse to select the block of cells that will serve as the named range. This is done by clicking the mouse in the upper-left cell of the block, and then, while holding down the shift key, clicking again in the lower-right corner of the block.



- From the Formulas tab, select “Name Manager” and then “New...”. This will open the Define Name dialog.



- Enter a name for the range. The name should be short (less than 30 characters long) and descriptive. It should begin with a letter, and contain only letters, digits and underscores. When the name has been entered, click the “OK” button to close the dialog.

9.3.1 Header Rows in Named Ranges

A named range must have a header row, if the Excel file is being used as a Dynamic Data Source (as opposed to being imported as a Static Data Source). A named range is said to have a header row when the top row contains column names rather than data (sometimes referred to as the first row). If a named range does not have a header row, iDashboards will

arbitrarily assign names to them when during chart creation (example: “column1”, “column2”, etc.). If a named range has a column header, and blank values are found, then names will be assigned using “F1” for the first blank field, “F2”, for the second blank field, etc.

The selected range has a header row, containing the column names “TheMonth”, “Sales”, and “New Customers”. When a header row is used, the column names are used as names for the columns of the virtual database table the named range represents. Column names should be less than 30 characters long, should begin with a letter and should only contain letters, digits and spaces.

9.4 Imported Files

Users with the role of “Data Admin” or “Admin” can import files. Using the “Select Files to Upload” button, choose the file or files to import.

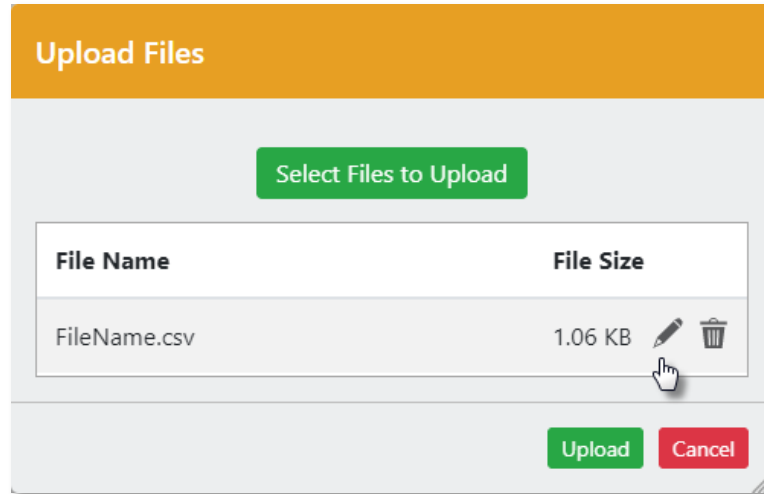
Single File

- 1) From ‘Imported Data menu, select either “Excel Data” or “Imported Files”
- 2) Click the button “Select Files to Upload”
 - a) Navigate to a single file and click “Open”
- 3) Click “Upload”

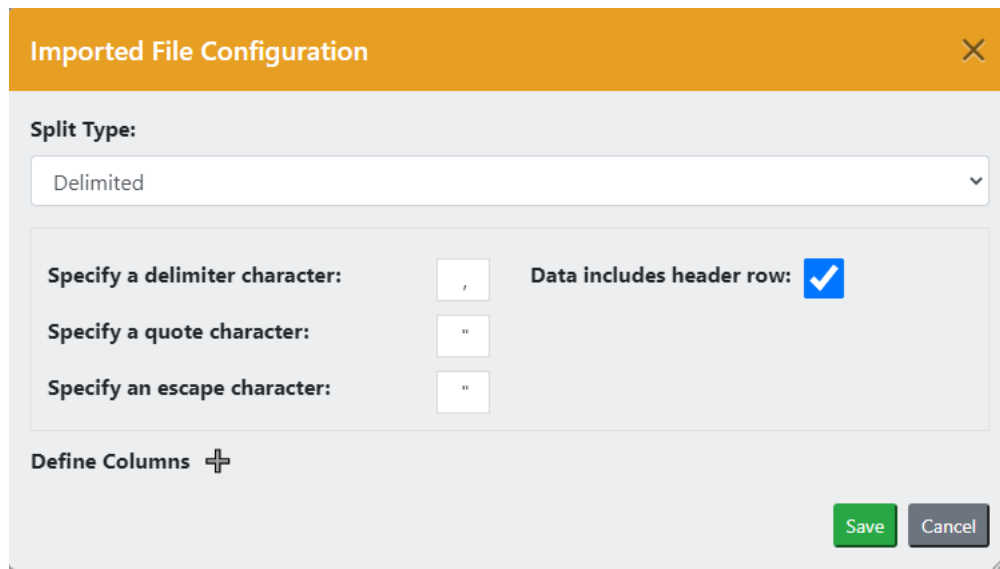
Multiple Files

- 1) From ‘Imported Data menu, select either “Excel Data” or “Imported Files”
- 2) Click the button “Select Files to Upload”
 - a) All in one folder
 - i) Select multiple files by holding down the CTRL key
 - ii) Click “Open”
 - b) Across multiple folders
 - i) Files from multiple folders need to be selected in batches of one folder at a time.
 - ii) Navigate to the first folder and select the file(s), then click “Open”
 - iii) Next, re-open the dialog by clicking on button “Select Files to Upload”.
 - iv) Navigate to the next folder and select the file(s), then click “Open”
 - v) Repeat
- 3) Click “Upload”

The primary difference between Excel Data and Imported File is that the administrator can specify how the Imported File is parsed. This is accomplished through the Imported File Configuration screen. It is access by using the pencil icon for the specific file.



For Delimited Files, user can define the columns, either in the file as a header record, or using the Imported File Configuration. Using this, a delimiter is specify and whether there is a header row. The default file extension is .csv. This is the default configuration:



If no columns are defined, the columns names and data types are inferred from the data. This is based on whether there is a header row and the `ibdata.text.max.rows.to.scan` setting in the `ivizgroup` property file. The default for this property is -1, which means scan all rows to determine the data type. If columns are defined, they take priority over inferred columns.

For Fixed Position files, all column must be defined using the Imported File Configuration. The default file extension is .txt. This is the default configuration:

Imported File Configuration [X]

Split Type:
Fixed Position [v]

Define Columns:

Column Name	Data Type	Start	End
Column1	String		

+ [trash] [Save] [Cancel]

For Regular Expression files, using the Imported File Configuration, a regular expression is provided along with capture groups. The Regular Expression pattern does not need to specify the entire string, but it does need to specify a substring, by using non-capture groups if necessary.

If not provided, the column names will be Column1, Column2, Column3, etc. Data types will be inferred from the data based on the `ibdata.text.max.rows.to.scan` setting, in the `ivizgroup` property file. The default for this property is -1, which means scan all rows to determine the data type. If column definitions are provided, the number of columns must match the number of capture groups. This is the default configuration:

Imported File Configuration [X]

Split Type:
Regular Expression [v]

Specify a Regular Expression:
[text input] [Test]

Ignore Case

Define Columns +

[Save] [Cancel]

9.4.1 File Name Rules

Under the hood, iDashboards requires uniqueness for the first 11-characters of a file. The first 11-characters are calculated by omitting any spaces and making all letters UPPERCASE.

For end-users, files without the first 11-characters being unique can be uploaded until there are 100 occurrences. When attempting to upload the 101st file (with a naming collision) the user will be notified of this rule and will have to change the filename before the upload will occur. The change needs to be a non-space character change within the first 11-characters.

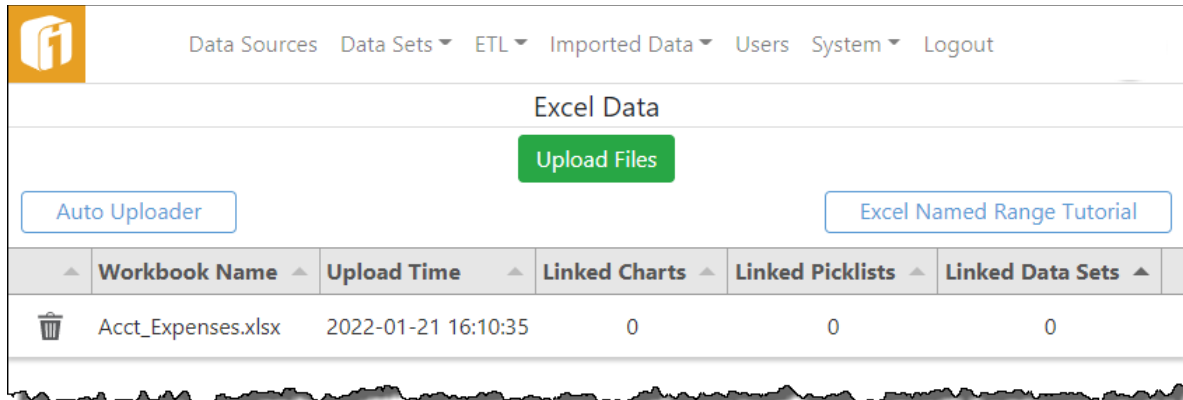
9.5 Automated File Uploads

To automate the process where files need to be uploaded, there is an automated file uploader. The iDashboards Auto Uploader allows for direct uploading of files from a client machine to the iDashboards Data Hub environment without requiring the user to log in to upload their data sources. The Auto Uploader allows files to be uploaded immediately or scheduled for repeated uploads using a scheduling mechanism.

9.5.1 Enable the setting

In the Data Hub, navigate to System > Settings and then review the option “Imported File Auto Upload Enabled”. This setting determines whether or not users can use the Auto Uploader desktop application to upload files to this server application. Set this to TRUE to enable the option.

Then, to access the iDashboards Auto Uploader and the associated documentation, log into the iDashboards Data Hub, navigate to “Excel Data” or “Imported Files” and click the button “Auto Uploader” and then follow the instructions.



9.5.2 Auto Uploader System and User Requirements

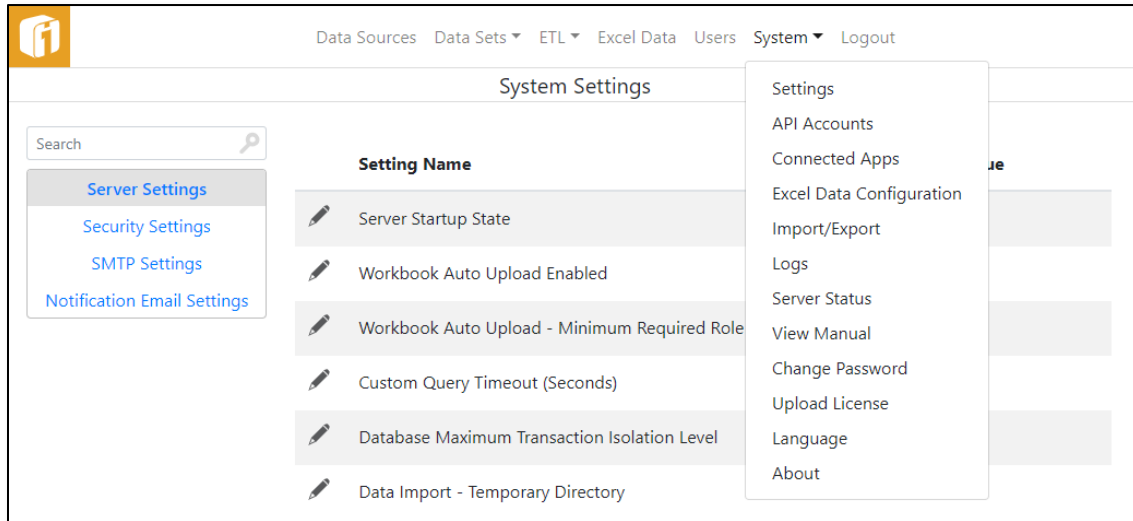
The IAU has the following minimum system requirements:

- Microsoft Windows 7 or Windows Server 2008 R2
- Microsoft .NET 4 Framework – Full profile


It also requires that the user running the program to have administrator privileges on the PC for which the iDashboards Auto Uploader is installed. These elevated privileges are necessary to create and execute tasks in the Windows operating system.

10. System

Various settings of the iDashboards Data Hub can be accessed by clicking “System” and then “Settings” from the application menu.



10.1 Modifying a System Setting

To modify a setting, click its Edit icon (). A form will appear through which the setting can be edited. For most system settings, the value must be entered into a textbox, while for some, the value can be selected from a dropdown list. The edit form for a system setting will include a description of that setting and its valid values.

After a system setting has been modified, click the “Save” button to save the changes.

10.2 Settings

“System settings” is divided into groups which correlate with the left-hand tabs. Available settings may differ from one version or installation of iDashboards to another, and may not match the list of system properties documented here. In a typical installation, system settings can be left at its default value. In some instances, however, a system setting might need to be changed to accommodate unique requirements.

10.2.1 Server Settings

Server Settings are global settings used to control the Data Hub server's behavior.

- **Server Startup State:** This setting determines the initial state of the server upon startup. The default value is "Running". The two possible values are:
 - **Running:** The Job Monitor Thread will be started, and the server will check for jobs to be executed according to their schedules.
 - **Paused:** The Job Monitor Thread will be in the paused state when the server starts up. It will need to be started manually through the Status section of the System tab in the Data Hub.
- **Imported File Auto Upload Enabled:** This setting determines whether or not users can use the Auto Uploader desktop application to upload workbooks and text files to this server application. The default value is "TRUE".
- **Imported File Auto Upload – Minimum Required Role:** This is the minimum role a user account must have in order to use the Auto Uploader desktop application. The default value is "Builder".
- **Custom Query Timeout (Seconds):** This setting indicates the maximum time, in seconds, that a custom query used to produce chart data will be allowed to run before it is terminated. Allowable values are from 1 to 120. The default value is 10. The default value is 10.
- **Database Maximum Transaction Isolation Level:** This setting should be left at its default setting unless you encounter error messages that say "Error setting isolation level" while doing things that modify the iDashboards repository database, such as adding or changing charts, dashboards, users, etc. The default value is Serializable.

If this type of error occurs, change this setting to reflect the maximum isolation level supported by your iDashboards repository database. (The levels increase toward the bottom of the list.) If you're not sure what the maximum supported isolation level is, use the "Database Default" setting.

- **Data Import - Temporary Directory:** When a user uploads an Excel or CSV file as imported chart data, iDashboards might need to temporarily write it to disk if it is unusually large. By default, it will write it to the system temp directory (c:\tomcat\Tomcat7.0\temp), but this setting can be used to indicate a different directory. The entered path can be absolute or relative to the appserver's working directory (c:\tomcat\Tomcat7.0), as long as it exists and is writable.

If iDashboards is running in a clustered appserver, the entered path must be valid for all nodes in the cluster. Therefore it's recommended to leave this setting blank in a clustered environment.

10.2.2 Security Settings

- **X-Frame-Options Header:** This property indicates to the browser whether or not the iDashboards Data Hub Application can be framed by another web page. Setting this property ensures the iDashboards Reports Application cannot be embedded into other web pages.

If the selected value of this property is “None”, then the X-Frame-Options response header will not be set, allowing the iDashboards Reports Application to be framed by another page. If the selected value is “Deny”, then iDashboards Reports Application can never be framed by another web page. If the selected value is “Same Origin”, then the iDashboards Reports Application can only be framed by a web page from the same domain. The default value is DENY.

- **Destroy Session on Login:** This property indicates whether or not an existing browser session should be destroyed when logging into the iDashboards Reports Application. The default value is TRUE.
- **Session Timeout (minutes):** Session Timeout (minutes).
- **Allow Sharing With External Applications:** Enabling this feature will allow data sets to be shared between Data Hubs or other applications. You will be able to configure individual sharing permissions on each data set. The default value is NO.
- **Require Secure Connection For Sharing:** If this setting is “Yes”, then the connection associated with the requests for shared data sets must be secure. If this setting is “No”, then the connections associated with the request is not required to be secure. The default value is TRUE.

10.2.3 SMTP Settings

The iDashboards Data Hub server uses an external SMTP service to send emails. On the SMTP Settings screen, locate the following settings:

- **SMTP Host:** This is the hostname or IP address of the machine on which the SMTP service is running.
- **SMTP Port:** This is the number of the TCP/IP port on which the SMTP service is listening. (The standard SMTP port number is 25.)
- **SMTP Service Requires Authentication:** Set this to “Yes” if the SMTP service requires authentication or incoming connections, or to “No” if it does not.
- **SMTP Service User:** If the SMTP service requires authentication, this setting must contain the username of the user that will be used to connect to it, otherwise it should be left blank.
- **SMTP Service Password:** If the SMTP service requires authentication, this setting must contain the password that will be used to connect to it, otherwise it should be left blank.

- **SMTP Encryption:** This setting determines the type of encryption (if any) used to secure the connection with the remote mail server. The options are “None”, “SSL (Secure Socket Layer)” and “TLS (Transport Layer Security)”.

10.2.4 Notification Email Settings

Notification emails are sent through an external SMTP service, such as UNIX Sendmail or Microsoft Exchange Server. The server must be configured with enough information to connect to, and if necessary, authenticate itself to the SMTP service.

- **Notification Email Enabled:** If this setting is “No”, then all email notifications, for jobs and server event notifications, will be disabled. If it is “Yes”, then the settings in the SMTP Settings category must be properly configured to connect to the SMTP Service. The default value is NO.
- **Other Settings:** Other notification email settings include information such as the name and email address used in the “from” header of outgoing emails, the list of email addresses that will receive server event notifications, and the information that is included in the subject lines of notification emails.

Configure the Email Templates: This is an optional process that provides a great deal of control over the information included in the bodies of notification emails. Using email templates, notification emails can be sent in both HTML format (including images) and plain text. If this step is omitted, emails will be sent as plain text and include only a minimal amount of default information.

10.3 Authentication Settings

Built into the core of iDashboards, is a framework to handle user authentication. Additionally, a system administrator can configure alternative authentication methods to assist with system integration or login automation needs. Refer to the iDashboards Administrator’s Manual for details.

10.4 API Accounts

To access API Accounts click “System” > “API Accounts”. This is also where the in-product documentation and instructions reside.

10.5 Connected Apps

Connected Apps is used for connecting one installation of the iDashboards Data Hub with other installations of the Data Hub for the purpose of sharing Data Sets. The Connected Apps page manages what external applications connect to the Data Hub. A Connected App contains all of the information required for an external application to connect using OAuth2.

To access Connected Apps, click “System” > “Connected Apps”.


To access the in-product documentation and instructions, click “System” > “API Accounts” and expand the button “Connected Apps”.

10.6 Imported Data Configuration

6) See Section 9.2, “Configure the Imported Data Database”

10.7 Importing and Exporting Data Sets and ETL Jobs

Data Sets and ETL Jobs can be exported from an iDashboards repository (called the source repository) and imported into another iDashboards repository (called the target repository). Data Sets and ETL Jobs are exported to an iDashboards archive file (named with the .zip extension). This file can then be referenced during the import process to import the Data Sets and ETL Jobs into the target iDashboards repository.



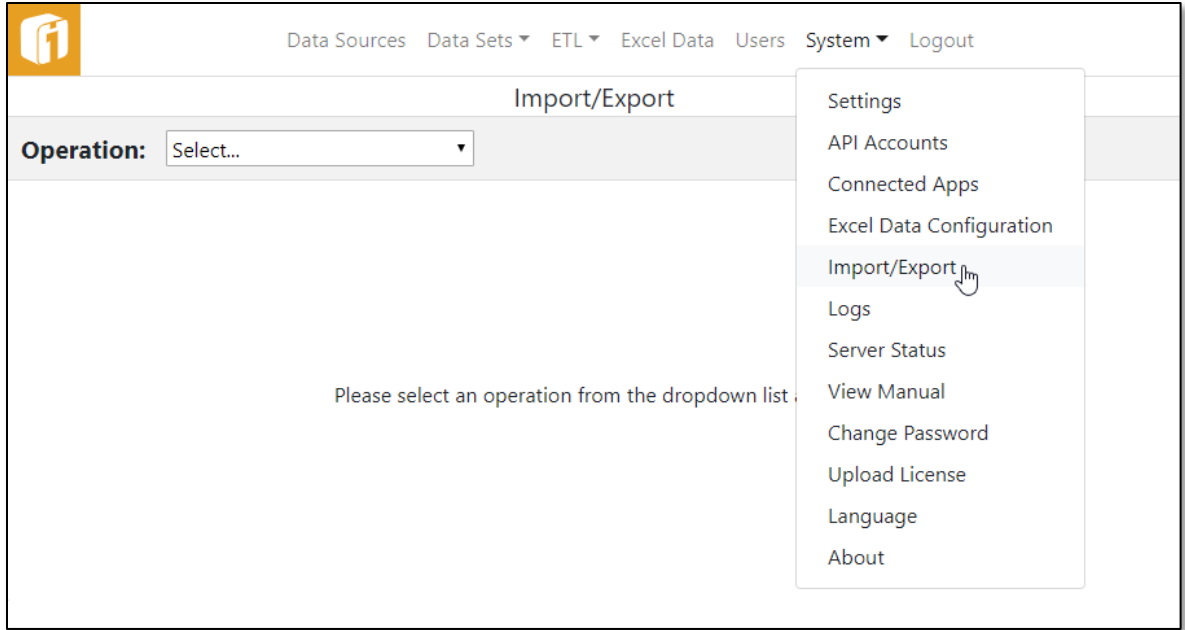
To fully support the import/export functionality, each data source, stored procedure and API account that is used by the Data Sets and ETL Jobs that will be exported and/or imported, must be assigned a Global Identifier (GID). GIDs are enterprise-wide identifiers that are used to identify data sources, stored procedures and API accounts across multiple installations of iDashboards. GIDs are not used for the Data Sets and ETL Jobs.

Note: The GIDs that are assigned to stored procedures are assigned to the stored procedure that is created within iDashboards, not the actual stored procedure that resides in the data source. The stored procedure that is created within iDashboards is, in essence, configuration information about the actual stored procedure that exists in the data source.

By default, GIDs **are not** assigned to data source, stored procedure and API account when they are created. An iDashboards administrator must manage GIDs as described later in this section.

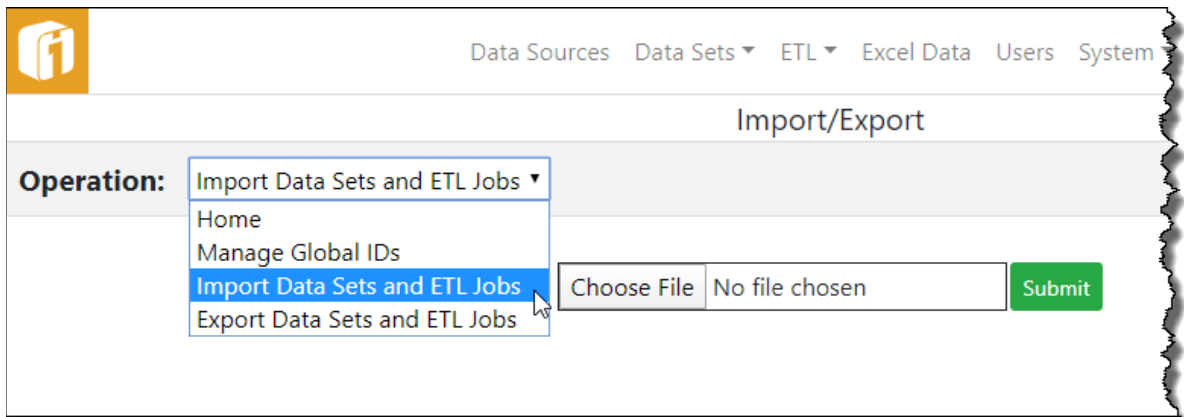
Note: GIDs are not the same as the local IDs that are automatically assigned.

The importing and exporting of ETL Jobs is managed through the “Import/Export” option under the System menu of the Data Hub Application.



The administrator has the option to select one of the following from the “Operation” dropdown list:

- Manage Global IDs
- Import Data Sets and ETL Jobs
- Export Data Sets and ETL Jobs



10.7.1 Criteria for Export

Data Sets and ETL Jobs may be exported separately. Each has their own criteria for exportation.

For a Data Set to be eligible for export, the following conditions must be met:

1. The Data Set's Data Source must be assigned a GID.
2. If that Data Source is a stored procedure, the stored procedure must be assigned a GID.
3. If that Data Source is an API Account, the API Account must be assigned a GID.

For an ETL Jobs to be eligible for export, the following condition must be met:

1. All Data Sets used in the ETL Job must adhere to those export requirements (described above).

10.7.2 Criteria for Import

Before importing takes place the data sources, and their Global ID's, must match between the two systems. The "IDB System Database" is the System Database which iDashboards uses to store metadata. By default the ID of this repository will be zero. However, if during the installation process the ID was purposely changed to another value, the ID will need to be changed within the second system before import. Other data sources must also have matching ID's before importing.

For a Data Set to be eligible for import, the following conditions must be met:

1. There must be a Data Source in the target repository with the same GID as the one associated with the Data Set that will be imported.
2. If a stored procedure is used to retrieve its data, the database the Data Set will point to must have a stored procedure identical to the stored procedure the imported Data Set pointed to in the source database.
3. If an API Account is used to retrieve its data, the database the imported Data Set will point to must have a API Account identical to the API Account the imported Data Set pointed to in the source database.

Note: Stored Procedure definitions in iDashboards don't need to be created manually in the target repository. They will be imported automatically when the .zip file is imported, assuming the stored procedure was given a GID during the export process. The stored procedure definitions are imported along with the charts that use them.

For an ETL Job to be eligible for import, the following condition must be met:

1. All Data Sets used in the ETL Job must adhere to those import requirements (described above).


10.7.3 Managing Global Identifiers (GIDs)

When an administrator selects "Manage Global IDs" from the "Operation" dropdown box, they will be presented with a list of the three iDashboards entity type groups that can be assigned GIDs (Data Sources, Procedures and API Account). GIDs are managed by selecting the group.

The screenshot shows the 'Import/Export' interface with the 'Operation' set to 'Manage Global IDs'. The main content area displays a table of data sources:

ID	Data Source Name	Global ID
7	AdventureWorks2014	
8	Data Store	
10	GeoPlot Data	
0	IDB System Database	
9	WORKBOOK DATABASE	

On the left sidebar, there are three main categories: 'Data Sources' (selected), 'Procedures', and 'API Accounts'. Below these, there are two buttons: a green arrow pointing to a hash symbol (#) and a red minus sign (-).

By default, no entities in iDashboards are assigned a GID when they are created. To change this, either select the  button on the “Import/Export” screen or set the system parameter called “Generate Global IDs” to True, under “Server Settings”, on the “System Settings” screen in the iDashboards application. When iDashboards is configured to generate GIDs automatically, every new entity that is created will simply get assigned the next numeric GID available. This is the reason GIDs are not assigned by default. For example, if two installations of iDashboards are configured to assign GIDs automatically, then the administrator has no control over the GIDs that are assigned to the various entities. Therefore, when the administrator chooses to import and export, they cannot be assured the GIDs will match up how they would like between the two installations. It is safer for the administrator to manually manage GIDs.

Selecting one of the groups (Data Sources, Procedures and API Account) will display a list of entities of the indicated type that are available for GID management. The list of entities will be grouped if applicable. Data Sources will not have any groupings. Procedures will be grouped by the Data Source they are associated with. API Accounts will be grouped by their API account type. To set the Global ID for those entities that are within a group you will need to click on their corresponding group to expand and view the entities. The expandable link will have a finger similar to a hyperlink when you hover over it.

The screenshot shows a web application interface for 'Import/Export'. At the top, there is a navigation bar with links: Data Sources, Data Sets, ETL, Excel Data, Users, System, and Logout. Below this is a header for 'Import/Export' and a dropdown menu for 'Operation' set to 'Manage Global IDs'. On the left, there is a sidebar with three categories: 'Data Sources', 'Procedures', and 'API Accounts'. The 'Procedures' category is selected, and a table is displayed with the following data:

ID	Procedure Name	Global ID
1	uspGetAddress	

At the bottom left, there are two icons: a green arrow pointing to a hash symbol (#) and a red minus sign (-).

If an entity has an associated Global ID it will be displayed in the “Global ID” column. To change or add a Global ID, enter the unique number in the “Global ID” column. The administrator must assign/update each GID individually.

Note: GIDs are assumed to be unique within an entity type; in effect, the same GID cannot be assigned to two different Data Sources, but a Data Source and a Procedure or API Account can have the same GID. If there is a GID conflict within an entity type, an error message will appear indicating that GIDs must be unique.

10.7.4 Exporting Data Sets and ETL Jobs

When an administrator selects “Export Data Sets and ETL Jobs”, from the “Operation” dropdown box, the export groups (ETL Jobs and Data Sets) will be available to select from. Selecting one of the groups will show a list of its available entities. To select the entity within the category you will need to check the box next to the corresponding item. An administrator will also have the option of selecting, or unselecting, all items in the category by checking the box at the top in the categories column header line.

The screenshot shows the 'Import/Export' interface with the 'Operation' set to 'Manage Global IDs'. The left sidebar has 'Data Sources' selected. The main area displays a table for 'Data Source: AdventureWorks2014'.

ID	Procedure Name	Global ID
1	uspGetAddress	

Below the table are two buttons: a green arrow with a hash symbol and a red minus sign.

Items that do not meet the criteria set forth in Section 13.7.1, “Criteria for Export” cannot be exported and are identified with a strikethrough on its name. If select a **✗** will show in the Status column. Clicking on the **✗** will bring up a dialog box that will describe why the entity cannot be exported (see Figure 10-8).

The screenshot shows the 'Import/Export' interface with the 'Operation' set to 'Export Data Sets and ETL Jobs'. The left sidebar has 'Data Sets' selected. The main area displays a table for '/Data Sets'.

Status	ID	□	Data Set
✗	2	<input type="checkbox"/>	Person
	3	<input type="checkbox"/>	Sales

A dialog box is open over the 'Person' row, containing the following text:

The data set "Person" cannot be exported for the following reasons:

- The data source "AdventureWorks2014" does not have a global ID assigned.

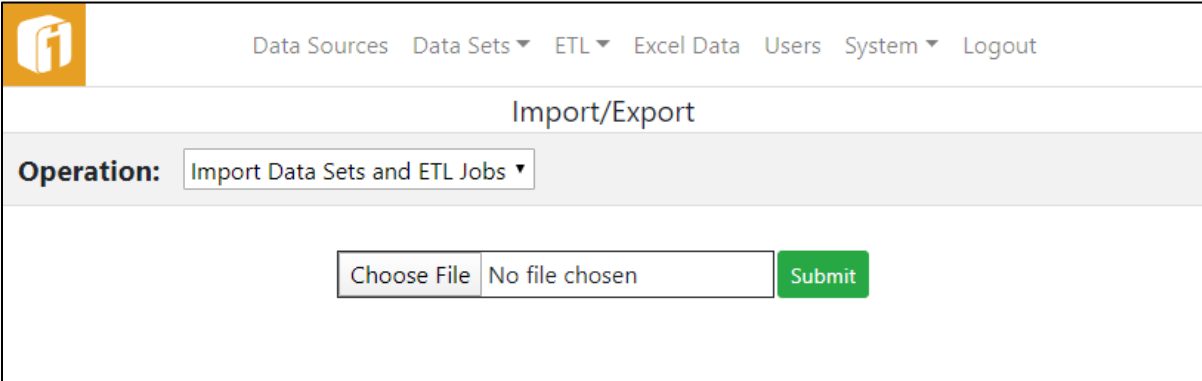
Buttons for 'Cancel' and 'Export' are visible at the top right of the interface.

There is the option to go back and cancel the whole export process by hitting the “Cancel” button. Once you have selected what is to be exported, click the “Export” button.

An iDashboards archive file is then created with the exported charts, dashboards and picklists. By default, this archive is named *etlarchive_yyyymmddhhmmssSSS.zip*. The name of the archive file may also be changed after the file is created and saved, but do not change the .zip file extension.


10.7.5 Importing Data Sets and ETL Jobs

When an administrator selects “Import Data Sets and ETL Jobs” from the “Operation” dropdown box, they will see a “Choose File” button and a “Submit” button. The “Choose File” button will assist in finding and filling out the text box which is the path of the .zip file. The administrator must do this on the iDashboards implementation that contains the target repository.



The screenshot shows the iDashboards interface for importing data. The top navigation bar contains links for Data Sources, Data Sets, ETL, Excel Data, Users, System, and Logout. The main heading is "Import/Export". Below this, the "Operation:" dropdown menu is set to "Import Data Sets and ETL Jobs". At the bottom of the form, there is a "Choose File" button, a text box containing "No file chosen", and a green "Submit" button.

The import functionality will create all of the entities from the archive file (Data Sets, ETL Jobs). The import function **will not** create any Data Sources nor assign any Data Source GIDs because any Data Source that an imported chart will use should already be created in the target repository (see Section 10.6.2, “Criteria for Import”).

A list of the Data Sets and ETL Jobs in the archive to be imported will be displayed. Select the  icon to view information about the import file.

The screenshot shows the 'Import/Export' interface. At the top, there are navigation links: Data Sources, Data Sets, ETL, Excel Data, Users, System, and Logout. The main heading is 'Import/Export'. Below this, the 'Operation' is set to 'Import Data Sets and ETL Jobs' and the file name is 'etlarchive_20200130103945158.zip'. There are 'Cancel' and 'Import' buttons. On the left, under 'Import Options', two checkboxes are checked: 'Import ETL jobs and data sets with warnings' and 'Overwrite existing ETL jobs and data sets'. On the right, there are two tables. The first table, titled '2 Data Sets', has columns 'Folder', 'Data Set Name', and 'Status'. It contains two rows: one for 'Person' and one for 'Sales', both with a green checkmark and 'New' status. The second table, titled '1 Etl Jobs', has columns 'Folder', 'Job Name', and 'Status'. It contains one row for 'ETL Job Alpha' with a green checkmark and 'New' status.

Data Sets and ETL Jobs will be validated before import. The status of the validation will be displayed on the staging window. If an item is valid for import then a ✓ will show next to the status on the staging window. If there are any problems importing an item, there will be a ✗ in the status column. Select the ✗ to see a description of why it cannot be imported. To allow importing items with warnings, under “Import Options” select the box “Import ETL jobs and data sets with warnings”.

This is a close-up of the 'Import Options' section. It shows two checkboxes. The first checkbox, labeled 'Import ETL jobs and data sets with warnings', is checked and circled in red. The second checkbox, labeled 'Overwrite existing ETL jobs and data sets', is unchecked.

After hitting the “Import” button a confirmation message is displayed notifying you of the import into the target repository.

10.7.5.1 Importing with existing GIDs

Caution should be taken whenever importing items into a non-pristine environment.

For example, a Data Set is imported with a given folder and name. If that folder and name already exists, then the import screen will display “Overwrite” in the status column. Continuing the import **will replace the existing data set with the imported data set**. If the

administrator does not want to import items with existing folder and name, but would like to import all other items, then, under “Import Options” unselect the box “Overwrite existing ETL jobs and data sets”.

Import Options

Import ETL jobs and data sets with warnings

Overwrite existing ETL jobs and data sets

Similarly, when an ETL Job is imported with a given folder and name, if that item exists with that folder and name, the existing ETL Job within the target repository will be replaced with the imported item. The folder and name of the dashboard or picklist is retained. If the item does not exist, then a new ETL Job with that folder and name is created in the target repository.

10.8 Log Settings

The Log settings screen can be accessed by clicking “System” then “Logs” from the navigation menu.

Once the server has been started, the settings can be changed through the Log Settings screen, however, changes made will not persist beyond a server restart. Under normal operating circumstances, the settings below should be used.

The screenshot shows the 'Logs' configuration interface. On the left, the 'Log Settings' section includes a dropdown for 'General Level' (set to 'Debug (most verbose)'), checkboxes for 'JSON Logging Enabled' and 'HTTP Request Logging Enabled' (both checked), and dropdowns for 'Max File Size' (100 MB) and 'Max Backup Files' (10). 'Reset' and 'Save' buttons are at the bottom. On the right, the 'Download Log Files' section has a list of log files ('idbdata.log', 'idbdata.log.1'), a 'Download...' button, and an 'Open Log Viewer' button.

Changes made to log settings are not applied until the “Apply” button has been clicked. The available settings are:

General Level — This setting determines the types of log messages that will be written to the log file. Each level can be thought of as a threshold, with Debug being the lowest and Error the highest. When a level is selected, all messages categorized at that level and above will be written to the log. The available levels are as follows:

- **Debug** — This is the most verbose setting and could impact system performance on a busy server. Debug log messages are generally only useful to an iDashboards support representative, so this level should only be used when troubleshooting.
- **Info** — This is a far less verbose level than Debug, which writes information about the operating environment to the log when the iDashboards server is started. **It is the recommended level for normal operations.**
- **Warn** — In addition to error messages, this level will write warning messages about server events that are noteworthy but not critical.

- **Error** — This is the least verbose log level. It will only write messages to the log when a critical error occurs.

JSON Logging Enabled – When this checkbox is checked, information about JSON requests that are sent to the iDashboards server are logged. When enabled, the logging is very verbose which is only useful to an iDashboards support representative, so it should remain unchecked except for troubleshooting purposes.

HTTP Request Logging Enabled — When this checkbox is checked, information about the HTTP requests that are sent to the iDashboards server are logged. As is the case with XML logging, it causes very verbose output which is only useful to an iDashboards support representative, so it should remain unchecked except for troubleshooting purposes.

Max File Size — This is the maximum size to which a log file will be allowed to grow before it is overwritten by a new one or archived.

Max Backup Files — This setting indicates the maximum number of archived log files that will be kept. When an active log file, named “idbdata.log” grows to its maximum allowed size, it will be renamed with a numeric suffix “idbdata.log.1” and a new idbdata.log file will be started. If there is already a file named idbdata.log.1 in the logs directory, it will be renamed idbdata.log.2, and so on, up to the maximum number of archived log files. When the maximum number has been reached, the oldest archived log file will be discarded as illustrated below.

Log File Name	Relative Date	Relative Age	Status
idbdata.log	Today	Newest	Active
idbdata.log.1	Last Week		Archived
idbdata.log.2	Last Month	Oldest	Archived

10.8.1 Downloading Log Files

The active log file (idbdata.log) and any existing archived log files (idbdata.log.1, idbdata.log.2, etc.) can be downloaded through the Log Settings screen. To do so, select the desired files from the list at the right of the screen and click the “Download...” button. The selected files will be bundled into a ZIP file by the server and downloaded.

10.8.2 Sending Log Files to iDashboards Technical Support

When working with iDashboards technical support to troubleshoot problems with the iDashboards server, it is useful to provide the iDashboards log file(s) to the support representative. Problems can be diagnosed and corrected more expeditiously if these steps are followed prior to contacting iDashboards technical support:

1. Set the General Level to Debug, and enable JSON logging and HTTP Request Logging.

2. Recreate the error condition through the User Application or the appropriate Administrator Application screen.
3. Download the idbdata.log file and the idbdata.log.1 file if it exists, as described earlier in this section.
4. Email the ZIP file containing the log file(s), along with a description of the problem (and steps to recreate it if possible) to support@idashboards.com.

10.8.3 Log Viewer

In addition to downloading and opening a log file, the Data Hub allows for a “near” real-time snapshot of the active log file. Click the button “Open Log Viewer” and an external browser window opens to display the last 2,500 entries from the active log file. Here, it is possible to perform live filtering by the type of entry (DEBUG, INFO, ERROR, and WARN) or by the request ID. (REQ-ID:).

The purpose of this interface is to quickly diagnose log entries while recreating a workflow using the main application. Below is a suggested workflow to best-utilize this feature when diagnosing a problem:

- 1) Open Log Viewer window
- 2) Return to the main window running the Data Hub Application
- 3) Execute a task that is not performing as expected
- 4) Switch to the Log Viewer window and click the “Refresh” button
- 5) Diagnose the latest entries
- 6) Copy & Paste those entries and send them to iDashboards Support

Note: The Log Viewer only displays the latest ~2,500 line items from the active log file. Archived log files are not queried. To diagnose log entries beyond the most current entries, refer to the entire documentation within the entire Section 10.6, “Log Settings”.

Note: The row numbers displayed in the preview window are not related to the actual line item number in the active log file. Clicking the “Refresh” button will completely rearrange the numbered index.

Note: The task of opening (and refreshing) the log viewer will generate logged items which display in the viewer.

10.8.4 Log Configuration

At runtime, iDashboards will log system errors and other events in a log file. The name of the log file is idbdata.log, and it will be created in the <IVIZGROUP HOME>\logs directory. Certain parameters can be set in the ivizgroup.properties file to determine the maximum size a log file will be allowed to grow to, the number of backups that will be kept, and the verbosity of the logging output. **Note that these settings can be changed while the server**

is running through the System Logs screen, described in Chapter 10, “System”, however such changes will not persist across a server restart.

log.directory – This property can be used to indicate a directory other than <IVIZGROUP HOME>\logs where log files should be written. It must exist and be writable by the iDashboards application server process. Forward slashes (/) should be used instead of backslashes (\) as a path separator.

log.maxFileSize – This property indicates the maximum size, in bytes, that a log file will grow to before it is “rolled over”, that is, renamed with a “.1” extension so that a new idbdata.log file can be created. This property must be an integer from 0 to 9,223,372,036,854,775,808. (Do not include commas.) The suffixes “KB”, “MB”, or “GB” can be appended to indicate the value is kilobytes, megabytes or gigabytes, respectively. If no value is given, the default used is “10MB”.

log.maxBackupIndex – When logs are rolled over, the current idbdata.log file is renamed to idbdata.log.1, an existing log file with a “.1” extension is renamed with a “.2” extension, one with a “.2” extension is renamed with a “.3” extension, and so on up to the value of the log.maxBackupIndex property. If a log file already has an extension equal to log.maxBackupIndex, it is discarded when the log files are rolled over. If log.maxBackupIndex is zero, there will be no backup files, and the log will be truncated when its size grows to the maximum size.

log.level – This value must be one of the following: ERROR, WARN, INFO or DEBUG. The default is INFO. DEBUG will produce the most verbose output, and ERROR will produce the least. Normally, DEBUG should only be used when troubleshooting.

10.9 Server Status

The Data Hub can be an active server depending on the count, complexity and frequency of running ETL jobs. The Server Status found beneath “Settings” is similar to the Job Status found beneath “ETL” – whereas the Server Status contains the status of all ETL jobs checked, executed and queued in addition to the overall run-state of the job engine. Here, administrators can monitor and troubleshoot the overall health of the Data Hub.

10.9.1 Pausing and Restarting the Server

At any given moment, the Data Hub Server will be in one of two possible states:

- **Running** – In this state, the Server is performing all of its normal activities, such as running scheduled ETL jobs, sending emails, etc.
- **Paused** – In this state, the Server does not perform activities such as running scheduled ETL jobs or sending emails, however, the Server console is still fully functional.

In its default configuration, the Data Hub Server enters the running state when it is started. When it is in the running state, the Server Status screen will display the line, “Current State: RUNNING”, and a toggle button will allow the administrator to “Pause Server”.

When the Data Hub Server is in the paused state, the Server Status screen will display “Current Status: PAUSED” and the toggle button will say “Start Server”. It can be placed back into the running state by clicking the button.

Normally, the Data Hub Server should be left in the running state. The paused state is generally only useful when performing troubleshooting or certain configuration changes.

10.9.2 Understanding Server Events

The most prominent feature of the Server Status screen is the list of server events. A server event can be any type of noteworthy occurrence, such as the server being paused or a database error. The event list can be filtered to only display events of certain, selected levels. This is accomplished by opening the search options and checking or unchecking the checkboxes for the different event levels.

A server event has the following attributes:

10.9.2.1 Event ID

Each server event is assigned a code referred to as the “event ID”, which identifies the type of event that it is. An event ID consists of an event category, such as “MONITOR”, and a number, separated by a hyphen.

The event category is used to identify approximately where in the system the event occurred. For example, the MONITOR category is for events that occur on the monitor thread, which is the main thread that runs continually inside the server, running scheduled ETL jobs and performing other tasks.

The number portion of the event ID uniquely identifies the type of event within an event category. For example, “MONITOR-7” is the event ID used to indicate that a routine job check occurred.

10.9.2.2 Level

Each server event has one of the following three levels:

- **INFO** – This level is used for routine events. INFO-level events are displayed in green text in the event list.
- **WARNING** – This level is for events that occur during normal operation, but should be noted by a server administrator. WARNING-level events are displayed in the yellow text in the event list.
- **ERROR** – This level is used for abnormal, unexpected events such as a database error that occurs during job generation. ERROR-level events are displayed in red text in the event list.

10.9.2.3 Timestamp

The event timestamp is the date and time at which the event occurred.

10.9.2.4 Subject

The event subject is a short phrase describing the event.

10.9.2.5 Message

The event message is a short sentence that contains information about the event.

10.9.3 Troubleshooting Error Events

In some cases, an ERROR-level event displayed in the event list may contain hyperlinks to other screens that display additional information about the error or the associated alert.

For example, if the event's level code, "ERROR", is followed by "(!)" the exclamation mark can be clicked on to display the Java stacktrace that was generated by the error.

Although stacktraces appear undecipherable to almost anyone other than Java developers, they usually provide clues as to the cause of an error. For example, a stacktrace notice might read "Connection timed out", indicating that the Job Server was unable to connect to the SMTP service to send emails.

A stacktrace can also be copied and pasted into an email to support@idashboards.com to assist iDashboards support staff in troubleshooting errors.

When an error is associated with a specific alert, the event message in the list may be followed by "(View Job)". This is a hyperlink that can be clicked to open the administrative screen for the errant job, through which the job can be temporarily disabled or permanently deleted.

10.9.4 Event Retention

During normal operation, the Data Hub Server is frequently recording new events in the event list. Because of this, one would expect that over time, the event list would grow extremely large, yet it does not. This is because only a certain number of events with a given event ID are retained in the event list. This number is referred to as the "retention depth" for that event ID. When the number of events with a particular event ID exceeds the retention depth for that ID, the oldest ones are removed from the list and discarded, keeping the entire event list at a manageable size.

The retention depth for an event ID is normally not of concern to the Alerts Server administrator. It can be viewed, however, by holding the mouse cursor over the event ID in the event's list. This will produce a tool tip, similar to the one shown below, displaying the retention depth for the event ID.

MONITOR-9	ERROR (!)	2014-08-15 17:00
Retention Depth: 1*		2014-08-15 15:33
Email: true		
STARTUP-2	INFO	2014-08-15 08:52

10.9.4.1 Qualified Event Retention

For some error events, the retention depth is not applied to the event ID alone, but rather to the event ID combined with some hidden qualifying information. For example, if the error event is related to a particular alert, that alert's ID number might be used as the qualifying information. So, if the event ID is "MONITOR-9" and the Job ID is 123, the hidden, "qualified" event ID to which the retention depth would apply would effectively (if not actually) be "MONITOR-9-123". And if the retention depth for MONITOR-9 events is 1, that really means that one MONITOR-9 event related to Job #123 will be retained in the list, but at the same time a MONITOR-9 related to Job #905 might be retained in the list as well. This keeps important events from being pushed out of the event list before they can be viewed by an administrator.

If a retention depth applies to a qualified event ID as described above, it will be followed by an asterisk (*) in the event ID tool tip.

10.9.5 Email Events

Certain event types are designated as "email events". When an email event occurs, a notification email will be sent to the designated Data Hub Server administrators, provided that:

- The Data Hub Server is properly configured to send event notification emails.
- The level of the event (INFO, WARNING, or ERROR) is at or above the configured threshold at which the event notification emails are sent.

To determine whether or not an event in the list is an email event, hold the mouse cursor over its event ID until the tool tip appears. It will include the line "Email: true" for email events, and "Email: false" for non-email events.

10.10 View Manual

Click this option to access this manual. Looks like somebody already did this!

10.11 Change Password

Clicking the "Change Password" link opens a dialog that allows you to enter your old password and then set a new password. If you forget your password, ask your iDashboards Administrator to assign you a new temporary password that you may change after you log-in.

Note: You cannot change your iDashboards username through this window.

10.12 Upload License

The iDashboards main application and the Data Hub utilize the same license file. Therefore, the task of uploading a new or updated license file can take place in either interface. The license file obtained from iDashboards will have the extension *.lic and can be uploaded from any client who has access to this interface.

Within the Data Hub, navigate to Settings > Upload License. Click on the button “Choose File” (or “Browse”) and locate the license.

After iDashboards has located and read the license file, its location will be displayed near the bottom of the Data Hub license screen. During the upload process, the file will be renamed to “idashboards.lic” and will be placed in the <IVIZGROUP HOME> directory. Refer to the Administrators Manual for more details.

10.13 Changing Language

Each user may choose which language iDashboards uses throughout the framework. To change the language preference, select “System > Language” and then the desired language from the list.

Languages are added, removed and updated by the iDashboards Administrator, by means of downloading and installing a language pack. Language packs affect the text that is built in to the iDashboards Application. For example, the language packs affect the menu items, navigation, messaging, etc. Language packs will not affect any text that is pulled from a Data Source as it comes directly from the data source. The iDashboards manuals are written in US English.

Note: Contact iDashboards for availability of Language Packs, or check OSKAR to download current Language Packs (see Chapter 2, “OSKAR”).

10.14 About

Click this link to view the Data Hub license and deployment information, including expiration date, number of licensed users and installed drivers. For archive or support purposes, click the button “Copy Table Information” to then paste the information into an email or text file.