



# Release Notes (2.8)

Version 2.8  
2025-03-03

*Release Notes (2.8)*

PDF generated on 2025-03-03

InterSystems® Data Fabric Studio™ Version 2.8

Copyright © 2025 InterSystems Corporation

All rights reserved.

InterSystems®, HealthShare Care Community®, HealthShare Unified Care Record®, IntegratedML®, InterSystems Caché®, InterSystems Ensemble®, InterSystems HealthShare®, InterSystems IRIS®, and TrakCare are registered trademarks of InterSystems Corporation. HealthShare® CMS Solution Pack™, HealthShare® Health Connect Cloud™, InterSystems® Data Fabric Studio™, InterSystems IRIS for Health™, InterSystems Supply Chain Orchestrator™, and InterSystems TotalView™ For Asset Management are trademarks of InterSystems Corporation. TrakCare is a registered trademark in Australia and the European Union.

All other brand or product names used herein are trademarks or registered trademarks of their respective companies or organizations.

This document contains trade secret and confidential information which is the property of InterSystems Corporation, One Congress Street, Boston, MA 02114, or its affiliates, and is furnished for the sole purpose of the operation and maintenance of the products of InterSystems Corporation. No part of this publication is to be used for any other purpose, and this publication is not to be reproduced, copied, disclosed, transmitted, stored in a retrieval system or translated into any human or computer language, in any form, by any means, in whole or in part, without the express prior written consent of InterSystems Corporation.

The copying, use and disposition of this document and the software programs described herein is prohibited except to the limited extent set forth in the standard software license agreement(s) of InterSystems Corporation covering such programs and related documentation. InterSystems Corporation makes no representations and warranties concerning such software programs other than those set forth in such standard software license agreement(s). In addition, the liability of InterSystems Corporation for any losses or damages relating to or arising out of the use of such software programs is limited in the manner set forth in such standard software license agreement(s).

THE FOREGOING IS A GENERAL SUMMARY OF THE RESTRICTIONS AND LIMITATIONS IMPOSED BY INTERSYSTEMS CORPORATION ON THE USE OF, AND LIABILITY ARISING FROM, ITS COMPUTER SOFTWARE. FOR COMPLETE INFORMATION REFERENCE SHOULD BE MADE TO THE STANDARD SOFTWARE LICENSE AGREEMENT(S) OF INTERSYSTEMS CORPORATION, COPIES OF WHICH WILL BE MADE AVAILABLE UPON REQUEST.

InterSystems Corporation disclaims responsibility for errors which may appear in this document, and it reserves the right, in its sole discretion and without notice, to make substitutions and modifications in the products and practices described in this document.

For Support questions about any InterSystems products, contact:

**InterSystems Worldwide Response Center (WRC)**

Tel: +1-617-621-0700

Tel: +44 (0) 844 854 2917

Email: [support@InterSystems.com](mailto:support@InterSystems.com)

# Table of Contents

<b>Release Notes (2.8)</b> .....	<b>1</b>
1 Business Scheduler .....	1
1.1 Tasks Can Be in Groups .....	1
1.2 Abort Option Now Aborts Scheduled Resource .....	1
1.3 New Run Modes .....	1
1.4 Run Reports Moved to Business Scheduler .....	1
1.5 More Granular Time Reporting .....	2
1.6 WHEN() and WINDOW() Functions .....	2
2 Configuration Management .....	2
3 Connectors .....	3
3.1 FileDir Connector Change .....	3
3.2 SQL/JDBC Connector .....	3
4 Recipes .....	3
4.1 Faster Recipe Reset .....	3
4.2 New Validation Functions .....	3
4.3 Reconciliation Corrections .....	3
4.4 Promotion to Files on S3 Buckets .....	4



# Release Notes (2.8)

This page describes the major changes in InterSystems Data Fabric Studio™ 2.8, including bug fixes that may affect how you work.

In addition to these changes, the entire product has been reviewed and updated to improve usability and increase consistency.

## 1 Business Scheduler

### 1.1 Tasks Can Be in Groups

Tasks in the Business Scheduler can now be organized into [groups](#), which you can collapse or expand. A group can have zero or more scheduled tasks of any type. The groups enable you to visually manage a potentially large number of tasks. The groups cannot be scheduled; their only purpose is to visually organize the Business Scheduler Dashboard.

For each task group, the Business Scheduler summarizes the statuses of all the tasks in the group; by design, this is the most important status of the group. For instance, if all tasks in a group are Running, but one task is in Error, the group summary status is Error. This way, even when a group is collapsed, the Business Scheduler displays the status that most needs attention.

### 1.2 Abort Option Now Aborts Scheduled Resource

The **Abort** button in the Business Scheduler Task Detail page has been enhanced to not only abort a scheduled task that is in the Running Wait status, but also abort the scheduled resource if it is already running, if applicable. For instance, a recipe can be aborted when it is running while a snapshot can't, so the Abort button will be enabled or disabled for a running task depending on its type.

### 1.3 New Run Modes

The Business Scheduler now supports two new run modes:

- **Run After**—The task will be run only after another task has been successfully run.
- **Run After Dependency**—The task will be run every time the dependency is satisfied. In contrast to the previously existing **Run on Schedule** mode, a dependency expression is *required*.

For details and examples, see [Scheduling a Task](#).

### 1.4 Run Reports Moved to Business Scheduler

For recipes, the system automatically generates a report that details the results of the execution of the work. (Snapshots and cube builds do not yet have such reports.)

The Business Scheduler Task Detail page shows the history of all runs of a given scheduled task ([run history](#)). With this improvement, each row in the run history will also show a link to the run report, if the scheduled task is a recipe.

It is important to notice that it is no longer possible to reach the run report from the recipe. Such information is now centralized in the Business Scheduler module. To see the run report for a recipe, find its scheduled task in the Business Scheduler, open the task detail so that the task run history is displayed and click the **Run Report** link of a specific run.

## 1.5 More Granular Time Reporting

The [Business Scheduler](#) main page and task details pages both now provide greater detail on the time spent to run the task:

- **Task Start Time**—When exactly this task started running.
- **Task Wait Time**—The amount of time spent on waiting for dependencies such as waiting for an event to arrive or waiting for a previous run of the job to finish (concurrency).
- **Job Start Time**—When the job (recipe, snapshot, cube, etc.) actually started running.
- **Job Run Time**—How long did it take for the job to finish running.
- **Job End Time**—When did the job finish running.
- **Total Elapsed Time**—How long did the entire task take to run (including waiting time).

## 1.6 WHEN() and WINDOW() Functions

When scheduling a task, depending on the scheduling mode, you can provide a [dependency expression](#). See [Managing Task Dependencies](#).

As of this release, a dependency expression can use the new WHEN() function. When a task starts running and it has the WHEN(SomeTag) function in its dependency expression, the system will hold the task in the RunningWait status until a business event identified by SomeTag arrives. This business event must arrive/happen *after* the task reaches the RunningWait status. When such a business event happens, the task will then proceed to Running unless the dependency expression has other conditions that must be satisfied.

Also, the WINDOW() function is now documented. This function enables you to specify a window of time during which the other task has completed. This window of time can span days. For details, see [Functions Available in Dependency Expressions](#). The WINDOW() function was present in previous releases, so this function has also been added to the older documentation.

# 2 Configuration Management

This version of the product brings several enhancements and bug fixes to the [Configuration Management](#) module. Some of the most important ones are as follows:

- It is now possible to import a configuration such as a recipe that already exists in the system with a given name, but a different GUID. Before this fix, the system would refuse to allow this configuration to be loaded because although they have the same name, there is no way for the system to know if this is the right thing to do or not. Two users could have worked on a recipe with the same name on two different environments and one of these users is now trying to bring his/her work to the other environment. In order to prevent an accidental loss of work, the system would prevent the user from loading that configuration. Now, the system will warn the user about the name collision and allow the user to override the protection and load it anyway.
- It is now possible to import/export:
  - Entities from the Entity Master
  - Scheduled tasks

- Recipe groups have been added to the YAML format for recipe configurations.

## 3 Connectors

### 3.1 FileDir Connector Change

When importing the schema of a CSV file into the Data Catalog, the import process now recognizes quote characters in the column headers of the file. For example, say a CSV file had "City, State" as one of its column headers. While a comma is the delimiter of the file, this comma is contained within quotes and should be escaped. Prior to the fix, City and State were split into two separate fields in the schema definition. With the fix in place, the header of the CSV file is parsed with the quote character factored in, so the field appears as City\_State in the schema definition instead. The comma has been converted to an underscore, but this aspect is pre-existing and the agreed upon business logic.

### 3.2 SQL/JDBC Connector

A bug has been fixed where the product would trigger a BUILD INDEX on all staging tables of staging activities that were left in the Staging Error status when the user asked to retry a failed recipe. This caused a lot of unnecessary data processing. This issue had been previously corrected in the case of SQLQuickLoad (LOAD DATA). The same correction has now been applied when not using SQLQuickLoad.

## 4 Recipes

### 4.1 Faster Recipe Reset

Resetting of recipes is now faster because the truncate and drop table instructions occur outside a transaction, so they will not generate journal entries.

### 4.2 New Validation Functions

New [validation functions](#) are available—you can now validate whether a column does/does not match a glob pattern, an InterSystems IRIS pattern, or a regular expression.

### 4.3 Reconciliation Corrections

This version includes the following corrections to reconciliation activities:

- When the user defines a reconciliation primary key map, the product will now automatically create indices on the staging table that will include the fields mentioned on the primary key map so that joining the two tables will be efficient. (Before this fix, the product prevented you from editing the reconciliation primary key map until after the indices were built, which you may have had to initiate.)
- Reconciliation now correctly handles the case where columns specified in the key map contained null values. (Recall that in reconciliation, the key map describes how to join the two tables in order to compare the records that match.) With this fix, a record is now considered to fail reconciliation if:
  - The record is in the primary source and has all null values for the columns in the key map.

- The record is in the secondary source, has all null values for the columns in the key map, and the reconciliation is a full outer join (that is, the reconciliation activity uses the option **Report Missing Records from Primary Source**).

As a part of the fix, failures resulting from null key map columns are also accurately reflected in a recipe's reconciliation reports. Note that are still attempted for records that have some null values for columns in the key map. For example, table A is being reconciled against table B. If table A has the columns FirstName and LastName in the reconciliation's key map, and one of its records has a null FirstName but a LastName value of "Smith", then the reconciliation will try to join to a record in table B with <null>,"Smith" in its key map columns. If there is no match, then that record's reconciliation fails in the same way it did before.

## 4.4 Promotion to Files on S3 Buckets

In this version, it is possible to promote data to CSV files or other files on S3 buckets. That is, when you define a [promotion activity](#), that activity can promote data to files on S3 buckets.