

# Relativity Data Grid Guide

---

September 29, 2020

For the most recent version of this document, visit our [documentation website](#).

# Table of Contents

<b>1 Relativity Data Grid</b>	<b>3</b>
1.1 Getting started	3
1.2 Supported and unsupported functionality	5
<b>2 Data Grid Text Migration application</b>	<b>6</b>
2.1 Special considerations	6
2.1.1 Before running text migration	6
2.1.2 Running text migration	6
2.1.3 Supported and unsupported functionality	7
2.2 Installing the Data Grid Text Migration application	7
2.2.1 Installing the text migration agents	8
2.3 Running text migration	8
2.4 Running text migration	8
2.4.1 Identifying workspaces to migrate	8
2.4.2 Identifying long text fields to migrate	8
2.4.3 Creating a text migration job	8
2.4.4 Running the Breakage Report	9
2.4.5 Creating and/or activating a dtSearch index	10
2.4.6 Resolving saved searches and views	11
2.4.7 Running a text migration job	12
2.4.8 Viewing migration errors	13

# 1 Relativity Data Grid

Relativity Data Grid allows you to store, search, and analyze extracted text and audit data at massive scale. Data Grid enables faster workflows through continuous indexing and distributed search to gain deeper insight to your extracted text and audit data. The benefits of using the Data Grid data store include:


- Scalability and more efficient review workflows on case sizes with 10 million or more documents and billions of audits.
- More performant database maintenance, backup, and upgrades.
- A reduction in SQL server database sizes, leading to better performing SQL databases.
- Increased visibility into reviewer productivity and system performance.

When you ARM restore to RelativityOne, extracted text is automatically stored in Data Grid. New workspaces created from default Relativity templates now have text automatically stored in Data Grid. Data Grid uses Elasticsearch to store and search audits. In RelativityOne, audits are automatically stored in Elasticsearch.

This page provides a brief description of Data Grid terminology, functionality, and important considerations to take into account .

## 1.1 Getting started

Use the following workflow to get started with Data Grid.



## 1.2 Supported and unsupported functionality

To store text in Data Grid, you can use the same default file share as your natives and images, or you can designate a file share specifically Data Grid text. Once you enable a long text field's access to Data Grid, you can't disable it, so it's important to understand the benefits and limitations of storing text in Data Grid as compared to storing text in SQL.

**Note:** When you ARM restore to RelativityOne, extracted text is automatically stored in Data Grid.

Supported functionality	Currently unsupported functionality
<ul style="list-style-type: none"><li>■ Import/export through the Relativity Desktop Client</li><li>■ Viewer</li><li>■ Preview</li><li>■ OCR</li><li>■ dtSearch indexing and searching</li><li>■ Persistent highlight sets</li><li>■ Processing</li><li>■ Relativity Integration Points</li><li>■ Analytics</li><li>■ ARM</li></ul>	<ul style="list-style-type: none"><li>■ Keyword search</li><li>■ SQL queries to long text fields stored in Data Grid</li><li>■ Adding extracted long fields stored in Data Grid to layouts (including the Document panel)</li><li>■ RSAPI query</li><li>■ Pivot and Sort in the UI</li><li>■ Filtering in the Document list on extracted text</li><li>■ Mass operations:<ul style="list-style-type: none"><li>○ Edit</li><li>○ Replace</li><li>○ Tally/Sum/Average</li><li>○ Export to File</li></ul></li></ul>

## 2 Data Grid Text Migration application

The Data Grid Text Migration application is a Data Grid application you can use to migrate your long text fields from SQL to Data Grid.

---

**Note:** Once you start a migration job, the results are permanent. You can't delete a migration job once it has been started.

---

### 2.1 Special considerations

There are several considerations you should be aware of before and during text migration.

#### 2.1.1 Before running text migration

Before running a text migration job, review the following considerations:

- If your workspace has any custom solutions or third-party applications built on the long text field you're migrating, update them to use the Export API to read long text and metadata fields.
- If you have any custom data which relies on the long text field you're migrating being stored in the SQL Document table, contact [Relativity Support](#) to discuss alternative solutions before migration.
- Data Grid only supports the **IS SET** and **IS NOT SET** operators. If your workspace uses other operators to query on long text fields, you will need to perform a similar query using dtSearch.
- Mass operations are not supported with Data Grid. If you need to populate a long text field stored in Data Grid, perform an OCR and select the Data Grid field as the destination field.
- Once the text migration job completes, the only index search available for the migrated field is dtSearch. Keyword search is not supported. Data Grid supports all features of dtSearch. If you do not have an active dtSearch index in your workspace, you will need to build one. You can include a combination of Data Grid and SQL fields in your saved search. For more information, see [Running the Breakage Report on page 9](#).

For a list of supported and unsupported Data Grid text functionality, see [Supported and unsupported functionality on the next page](#).

#### 2.1.2 Running text migration

Review the following considerations for running a text migration job:

- The Data Grid Text Migration application runs at the instance level. However, you must install the application to at least one workspace.
- If you have index build in progress (like dtSearch or Analytics), let the index build finish before starting a migration job.
- If you have a mass replace in progress, let the operation finish before starting a migration job.
- We recommend running a migration job in off hours. However, it is not required.

- You can only have one job in progress at a time. If you start a new text migration job while another is in progress, the new job is added to the queue.
- Although the application automatically deletes the column in SQL, you are still required to reclaim the space that was taken up by your long text document field in SQL.

### 2.1.3 Supported and unsupported functionality


Once you enable a long text field's access to Data Grid, you can't disable it, so it's important to understand the benefits and limitations of storing text in Data Grid.

Supported extracted text functionality	Currently unsupported functionality
<ul style="list-style-type: none"> <li>■ Import/export through the Relativity Desktop Client</li> <li>■ Viewer</li> <li>■ Preview</li> <li>■ OCR</li> <li>■ dtSearch indexing and searching</li> <li>■ Persistent highlight sets</li> <li>■ Processing</li> <li>■ Relativity Integration Points</li> <li>■ Analytics</li> <li>■ ARM</li> </ul>	<ul style="list-style-type: none"> <li>■ Keyword search</li> <li>■ SQL queries to long text fields stored in Data Grid</li> <li>■ Adding extracted long fields stored in Data Grid to layouts (including the Document panel)</li> <li>■ RSAPI query</li> <li>■ Pivot and Sort in the UI</li> <li>■ Filtering in the Document list on extracted text</li> <li>■ Mass operations: <ul style="list-style-type: none"> <li>○ Edit</li> <li>○ Replace</li> <li>○ Tally/Sum/Average</li> <li>○ Export to File</li> </ul> </li> </ul>

## 2.2 Installing the Data Grid Text Migration application

The Data Grid Text Migration application runs at the instance level. However, you must install the application to at least one workspace.

To install the Data Grid Text Migration application:

1. Navigate to the **Application Library** tab.
2. Select **Data Grid Text Migration** from the list of Relativity applications.
3. Next to Workspaces Installed, click **Install**.
4. Next to Workspaces, click , and then select a workspace.
5. Click **Ok**.
6. Click **Save**.

Once you install the Data Grid Text Migration application, the Text Migration Jobs tab appears under the Data Grid tab.

## 2.2.1 Installing the text migration agents

After you install the Text Migration application, you must install one Data Grid Migration Manager agent and at least one Data Grid Migration Worker agent to your environment. We recommend installing two Data Grid Migration Worker agents.

---

**Note:** You can run the Breakage Report without installing the text migration agents. We recommend running the Breakage Report and resolving any issues before running a migration job.

---

For more information on installing agents, see the Agents guide.

## 2.3 Running text migration

Once you install the Data Grid Text Migration application and configure your migration agents, you can run text migration. See [Running text migration below](#).

## 2.4 Running text migration

To run a text migration job, complete the following steps:

1. [Identify workspaces to migrate](#)
2. [Identify long text fields to migrate](#)
3. [Create a text migration job](#)
4. [Run the breakage report](#)
5. [Create and/or activate a dtSearch index](#)
6. [Resolve saved searches and views](#)
7. [Run text migration](#)
8. [Review and resolve migration errors](#)

### 2.4.1 Identifying workspaces to migrate

Identify the workspace(s) you want to migrate. Be sure to review the special considerations for text migration as well as the list of supported and unsupported functionality for Data Grid. For more information, see [Special considerations on page 6](#).



### 2.4.2 Identifying long text fields to migrate

Identify which long text field(s) you want to migrate. You can migrate up to three fields per migration job.

### 2.4.3 Creating a text migration job

To create a text migration job, complete the following:



1. Navigate to the **Text Migration Jobs** tab underneath the Data Grid tab.
2. Click **New Text Migration Job**.
3. Complete the following fields
  - **Name** - enter a name for the migration job.
  - **Workspaces** - click , and then select the from the list on the left the workspaces you want to migrate. Use the arrows to move your selection(s) to the list on the right. Once you are finished, click **Save**.
  - **Fields** - click , and then select the from the list on the left the fields you want to migrate. This list contains all long text fields that aren't Data Grid enabled. Use the arrows to move your selection(s) to the list on the right. Once you are finished, click **Save**.

---

**Note:** You can migrate up to three fields per job.

---

  - **Email notification Recipients** - enter the email address(es) of the recipient(s) you want to send a notification to when your migration job finishes running. Separate entries with a semi-colon.
4. Click **Save**.

#### 2.4.4 Running the Breakage Report

To run the Breakage Report, click **Breakage Report** from the text migration console. The Breakage Report provides a list of all views, saved searches, indexes, and custom objects that will no longer function once you migrate SQL text to Data Grid. We recommend resolving these issues before running a migration job.

One of the primary functions of the Breakage Report is to identify areas where any of the objects above are tied to a keyword search or using keyword search specific functions like Contains or Does Not Contain. Most of the issues identified can be resolved by modifying the object to use an active dtSearch index.

The Breakage Report contains the following columns:

Breakage Report - Chris's Test Run

The Breakage Report displays Views, Saved Searches, and Custom Objects that will no longer function in the event that you run Chris's Test Run

Workspace Id	Workspace Name	Name	Object Type	Owner	Field
1281017	SQL Workspace - EDRM 500k	Broken 4	Search	Public	Alert
1280873	SQL Workspace 15k - 1	Broken Search 1.3	Search	Martin, Alex	Alert
1280876	SQL Workspace 15k - 4	Broken 4	Search	Public	Alert

- **Workspace ID** - the Artifact ID of the workspace.
- **Workspace Name** - the name of the workspace.
- **Name** - the name of the view, saved search, index, or custom object that will break upon migration.
- **Object Type** - search, view, index, or custom object.
- **Owner** - the owner of the view, saved search, index, or custom object that will break upon migration.
- **Field** - the name of the field causing the object to break.
- **Operator** – the search operator (is, is not, etc.) incompatible with Data Grid.
- **Term** – the text string used in the broken search.

## 2.4.5 Creating and/or activating a dtSearch index

Most of the issues identified by the breakage report can be resolved by modifying the object to use an active dtSearch index. To create or activate a dtSearch index, see the Searching Guide.

**Note:** Ensure all fields you want to search on are included in the searchable set for the dtSearch index. You can include a combination of Data Grid and SQL fields in your saved search.

The Data Grid Text Migration application includes logic to automatically fix active dtSearch indexes using the **<all documents in workspace>** searchable set. This logic only applies if the fields being migrated have the Include in Text Index field set to Yes.

The Data Grid Text Migration application creates a saved search called **SS\_TextMigration\_[JobName]** which contains all fields where Include in Text Index is set to Yes. It also creates a new dtSearch index called **Migration\_[existing dtSearch index name]** which uses this new saved search. Once text migration completes, navigate to the old dtSearch index and click **Swap Index**. Select **Migration\_[existing dtSearch index name]** as the replacement index.

## 2.4.6 Resolving saved searches and views

Use the results of the breakage report to help you resolve index or field-level search issues that affect saved searches and views

---

**Note:** We recommend copying saved searches prior to making edits.

---

### 2.4.6.1 Resolving index searches

1. Navigate to the saved search.
2. Right-click on the name, and then click **Edit**.
3. Click the (Index Search) condition.
4. Copy the terms in the **Search Terms** field.
5. Change the **Index** field to an active dtSearch index.
6. Paste the search terms if needed.
7. Click **Apply**.
8. Click **Save & Search**.

### 2.4.6.2 Resolving field-level searches

1. Navigate to the saved search.
2. Right-click on the name, and then click **Edit**.
3. Click the condition using the field being migrated (for example, extracted text).
4. Copy the text query, and note the operator (is, is like, contains).
5. Click **Add Condition**, and then select (Index Search).
6. Next to **Index**, select an active dtSearch index.
7. Paste the query text you copied in step 4. Depending on the operator that was used, you may need to convert your search. For more information, see [Search operator conversion below](#).
8. Click **Apply**.
9. Remove the field condition.
10. Click **Save & Search**.

#### Search operator conversion

Data Grid only supports the **IS SET** and **IS NOT SET** operators. If your field-level search uses another operator, you need to convert the search as follows:

Operator	Example search	Conversion
is	Jane has a broken search	"Jane has a broken search"

Operator	Example search	Conversion
is not	Jane has a broken search	NOT "Jane has a broken search"
is set	<i>Supported operator</i>	<i>No modification required</i>
is not set	<i>Supported operator</i>	<i>No modification required</i>
is less than	N/A	N/A
is greater than	N/A	N/A
is less than or equal to	N/A	N/A
is greater than or equal to	N/A	N/A
is like	Jane has a broken search	*Jane has a broken search*
is not like	Jane has a broken search	NOT *Jane has a broken search*
begins with	Jane has a broken search	Paul*
does not begin with	Jane has a broken search	NOT Paul*
ends with	Jane has a broken search	*broken search
does not end with	Jane has a broken search	NOT *broken search
contains	Jane has a broken search	*Jane has a broken search*
does not contain	Jane has a broken search	NOT *Jane has a broken search*

## 2.4.7 Running a text migration job

Once you've resolved or taken note of any items in the breakage report, click **Start Job** to start the text migration job. If another text migration job is already in progress, the new job is added to the queue and begins as soon as any running or pending jobs complete

The job displays one of the following statuses which you can use to monitor the state of your migration:

- Pending
- In Progress
- Complete
- Completed with Errors

The **Migration Job Details** table shows the progress of each workspace in the migration job. This table refreshes every two seconds. Once the job completes or completes with errors, you can filter the table.

The table contains the following columns:

- **Workspace Name** - the name of the workspace being migrated.
- **Fields** - the name of the field being migrated
- **Migration Status** - the status of the field being migrated. The following lists the field migration statuses:

- Pending
  - In Progress
  - Not In Workspace
  - Already Migrated
  - Completed
  - Completed with Errors
- **Documents Migrated** - the count of documents migrated
  - **Total Documents** - the total number of documents migrated.
  - **Errors** - the count of documents with errors.

After a successful migration, the application runs a verification step to ensure all documents migrated from SQL to Data Grid. Then, the application drops the SQL column, permanently deleting the migrated text from SQL.

---

**Notes:**

- If all documents are migrated, but a workspace and field are still marked as In Progress, this means they are currently being verified.
  - Although the application automatically deletes the column in SQL, you are still required to reclaim the space that was taken up by your long text document field in SQL.
- 

## 2.4.8 Viewing migration errors

Run the **Data Grid Migration Error Report** from the **Text Migration Reports** tab to view migration errors. To run this report, click **Run**. Click the drop-down to toggle between Field Level Errors and Document Level Errors. Once you've resolved any errors, return to the migration job and click **Retry**. The job status returns to In Progress and any errors on the fields marked with errors are reset.


User Status Workspaces User and Group Management Queue Management Applications & Scripts Data Grid

Ingestion Management Text Migration Jobs **Text Migration Reports** Data Grid Text Settings

### Migration Reports

● Data Grid Migration Error Report  
This Relativity script reports on all errors that have occurred while migrating SQL Long Text data into the Data Grid repository.

Preview Run

Document Level Errors  Items 1 - 13 (of 13) [Filter] [Refresh] [Grid]

#	Field Level Errors	Workspace Id	Workspace Name	Field Name	Document ID	Error Message
	Document Level Errors	(All)	(All)	(All)	(All)	(All)
1	Text Load test - Text Migration	1298214	EDRM SQL Ex Text - Text Migration Load Testing	Extracted Text	3319736	- DataGrid.Exceptions.DataGr Error writing documents to Data Grid file system Response Result: ",3c85b069-e90f-4f1f- 9ee3- 7d312d2be388','Error'," 'Fields.ExtractedText','Error', network path was not found. ' at DataGrid.Implementations.D <WriteRecords>d__6.MoveN --- End of stack trace from previous location where exception was thrown --- at System.Runtime.ExceptionS at System.Runtime.CompilerSe task) at DataGrid.Imolementations.Fi

All 13 Export to File Go Viewing all items in sets of 25 per page

#### 2.4.8.1 Field Level Errors

The following results populate the bottom of the window:

- **Job Name** - the name of the migration job where the errors occurred.
- **Workspace Id** - the artifact ID of the workspace.
- **Workspace Name** - the name of the workspace.
- **Field Id** - the Artifact ID of the field with errors.
- **Field Name** - the name of the field with errors.
- **Error Message** - the text of the error.

#### 2.4.8.2 Document Level Errors

The following results populate the bottom of the window:

- **Job Name** - the name of the migration job where the errors occurred.
- **Workspace Id** - the artifact ID of the workspace.
- **Workspace Name** - the name of the workspace.
- **Document ID** - the artifact ID of the document.

- **Field Name** - the name of the field with errors.
- **Error Message** - the text of the error.

## Proprietary Rights

This documentation (“**Documentation**”) and the software to which it relates (“**Software**”) belongs to Relativity ODA LLC and/or Relativity’s third party software vendors. Relativity grants written license agreements which contain restrictions. All parties accessing the Documentation or Software must: respect proprietary rights of Relativity and third parties; comply with your organization’s license agreement, including but not limited to license restrictions on use, copying, modifications, reverse engineering, and derivative products; and refrain from any misuse or misappropriation of this Documentation or Software in whole or in part. The Software and Documentation is protected by the **Copyright Act of 1976**, as amended, and the Software code is protected by the **Illinois Trade Secrets Act**. Violations can involve substantial civil liabilities, exemplary damages, and criminal penalties, including fines and possible imprisonment.

©2020. Relativity ODA LLC. All rights reserved. Relativity® is a registered trademark of Relativity ODA LLC.