

Searching Guide

July 11, 2025

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

Table of Contents

1 Searching	9
1.1 Filters	9
1.2 Search panel	9
1.3 Keyword searches	9
1.4 Saved searches	9
1.5 dtSearch	9
1.6 Analytics	10
1.7 Regular expressions	10
1.8 Additional features	10
2 Filters	11
2.1 Showing and hiding filters in the item list	11
2.2 Filter types	12
2.2.1 Using Boolean filters	13
2.2.2 Using numeric filters for numbers	13
2.2.3 Using List filters	14
2.2.4 Using date filters	16
2.2.5 Using multi-list filters	18
2.2.6 Using text box filters	21
2.2.7 Text box filter search examples	22
2.2.8 Using advanced text box filtering	24
2.3 Changing item sets per page	25
2.4 Saving filters as a search	25
3 Keyword search	26
3.1 Fields	26
3.2 Example keyword search strings	27
3.3 Using the NOT operator in keyword searches	28
3.4 Understanding noise words	28
3.4.1 Keyword search noise words - without double quotes	28
3.4.2 Keyword search noise words - with double quotes	28
3.4.3 Single digits as noise words	29
3.4.4 Punctuation as noise words	29
3.4.5 At sign (@) and dashes	29

3.4.6 Hyphens and dashes	29
3.4.7 Default noise word list	29
3.5 Running a keyword search	30
3.5.1 Running a keyword search in the search panel	30
3.5.2 Running a keyword search in the Search browser	30
4 Search panel	32
4.1 Permissions	32
4.2 Collapsing, expanding, and moving the search panel	32
4.3 Condition card icons	32
4.4 Using the auto-run search setting	33
4.5 Creating a search in the search panel	33
4.6 Applying logic groups to search conditions	34
4.7 Navigating the search panel	36
5 Searching workflows	39
5.1 Workflow details	39
5.2 Best practices for advanced operators	39
5.2.1 Proximity search	39
5.2.2 Auto-Recognition	41
5.2.3 Searching for times	41
5.2.4 Filters	42
5.2.5 Troubleshooting workflow	42
6 Search conditions	44
6.1 Setting up search conditions	44
6.1.1 Setting up search conditions in the Search panel	44
6.1.2 Setting up search conditions in the Search browser	44
6.1.3 Using the multiple object condition builder	45
6.2 Canceling queries	45
6.3 Frequently asked searching questions	46
6.3.1 Multiple terms in dtSearch queries	46
6.3.2 Proximity searches in dtSearch queries	46
6.3.3 Using reserved characters in dtSearch queries	46
6.3.4 Multiple conditions	46
6.3.5 Nesting searches	47
7 Multiple object searching	48

7.1 These Conditions	48
7.2 These Conditions and These Conditions	49
7.3 NOT These Conditions	50
7.4 These NOT	50
7.5 None of these and Not all of these operators	51
8 Searching with the Entity object and name normalization	53
8.1 Considerations	53
8.2 Emails within a specific organization	53
8.3 Communications between two specific individuals and no one else	54
8.4 Communications between any two individuals and no one else	58
8.5 Emails across entity metadata	58
8.6 Emails sent from two different internal domains	60
8.7 Emails someone sent to themselves	61
8.8 Entities that communicated on privileged documents	62
8.9 Emails where a specific person drops off of a communication	63
9 dtSearch	65
9.1 dtSearch index	65
9.1.1 Accent-insensitive indexes	65
9.1.2 Creating a dtSearch index	66
9.1.3 dtSearch index page	67
9.1.4 dtSearch console	70
9.1.5 Re-index documents with incremental builds	71
9.2 Making the dtSearch noise word list searchable	72
9.2.1 Default noise word list	72
9.2.2 dtSearches and noise words	73
9.2.3 Noise words in languages other than English	73
9.3 Alphabet list	73
9.3.1 dtSearch Alphabet File	73
9.3.2 Alphabet file validation	75
9.3.3 Alphabet file sections	76
9.3.4 Reserved characters in the alphabet file	78
9.3.5 Searching for symbols and emojis	79
9.3.6 Setting up CJK document workspaces in Relativity	83
9.4 Running a dtSearch	88

9.4.1 Considerations	88
9.4.2 Running a dtSearch in the search panel	89
9.4.3 Running a dtSearch in the Search Browser	90
9.4.4 Search string examples and expected results	91
9.4.5 Searching for words longer than 32 characters	92
9.4.6 Using dtSearch syntax options	93
9.5 Running a dictionary search	112
9.5.1 Running a dictionary search in the search panel	112
9.5.2 Running a dictionary search in the search browser	114
9.6 dtSearch queue admin	117
9.6.1 Permissions	117
9.6.2 Monitoring dtSearch indexing jobs	117
10 Searching with regular expressions (regex)	119
10.1 Use cases for regular expressions	119
10.2 Regular expression metacharacters	120
10.2.1 Regular expression quantifiers	122
10.2.2 Escaping regular expression metacharacters	123
10.3 Using regular expressions with dtSearch	123
10.3.1 Regular expression search strings	123
10.3.2 Regular expression metacharacters	124
10.3.3 Regular expression groups	125
10.3.4 Escaping regular expression metacharacters	126
10.3.5 Common dtSearch regular expression examples	126
11 Saved search	128
11.1 Required security permissions	128
11.2 Navigating the saved searches browser	128
11.2.1 Filtering the list of saved searches	130
11.2.2 Performing mass operations on saved searches	131
11.3 Controlling the visibility of saved searches	133
11.4 Organizing saved searches in folders	133
11.4.1 Adding sub-folders to the root	133
11.4.2 Managing subfolders	133
11.4.3 Adding existing searches to folders	133
11.5 Creating or editing a saved search	134

11.5.1 Considerations	134
11.5.2 Create a saved search	134
11.5.3 Search fields	135
11.5.4 Pop-up pickers	138
11.5.5 Search bar	139
11.5.6 Link a Relativity application to a saved search	140
11.5.7 Link a dashboard to a saved search	141
11.5.8 Rerun out-of-date saved searches	141
11.5.9 Required security permissions	141
11.6 Defining criteria for saved searches	142
11.6.1 Operators	142
11.6.2 Batch fields as search conditions	146
11.6.3 Using saved searches as conditions (combined searches)	147
11.6.4 Lists as search conditions	148
11.7 Saving searches on the Documents tab	148
11.8 Common Saved Searches application	149
11.8.1 Installing the application	149
11.8.2 Using the application	150
11.9 Saved search history	151
12 Search terms reports	153
12.1 Permissions	153
12.2 Create and edit search term reports	154
12.2.1 Considerations	154
12.2.2 Create search term report	155
12.2.3 Editing search term reports	159
12.3 Tag Hits	160
12.3.1 Tagged documents in the Field Tree browser	161
12.3.2 Tagged search terms in a saved search	162
12.4 Run a search terms report	163
12.5 Search terms report hit count	165
12.5.1 Prerequisites	165
12.5.2 Saved search	166
12.5.3 Import	166
12.5.4 Run script	167

12.6 Errors	168
12.6.1 Syntax errors	169
12.6.2 Support errors	169
13 Short message search (Advanced Access)	171
13.1 Considerations	171
13.2 Creating a short message search	171
13.3 Defining search criteria operators	173
13.3.1 Long text operators	173
13.3.2 Fixed-length text operators	175
13.3.3 Date operators	176
13.3.4 Whole number operators	177
13.3.5 Yes/No operators	177
13.4 Searchable fields	178
13.5 Short message search index (Advanced Access)	180
13.5.1 Installing the application	180
13.5.2 Short message search index workflow	181
13.5.3 Retry errors	183
14 Search Q&A (Advanced Access)	185
14.1 Release notes	185
14.2 Considerations	185
14.3 Installation	186
14.4 Prerequisite	186
14.5 Using Search Q&A	186
14.6 Deleting an index	189
14.7 Q&A panel	190
14.7.1 Working with responses	192
14.8 Troubleshooting	193
14.9 Removing the Search Q&A application	193
14.10 Frequently asked questions	193
15 Field categories	195
15.1 Creating a field category	195
15.2 Viewing fields associated with a field category	195
15.3 Link a field or fields to a field category	196
15.4 Unlink a field or fields from a field category	197

1 Searching

Relativity includes flexible search features designed to facilitate the document review process. These features support a range of searching needs from filtering on fields and simple keyword searches to the development of complex queries. The following list summarizes the searching features available in Relativity.

1.1 Filters

You can use filters to limit the documents or items that appear in item lists on Relativity tabs and pop-ups. When you enable the filters for an item list, you can set criteria on single or multiple fields so that only matching documents or items appear in the view.

Filters query across the searchable set of documents in the active view to return your results. Relativity supports multiple filter types so that you can choose the best format for different field types. See the Searching Guide for more information.

1.2 Search panel

The search panel is available for the Document list and for many other tabs in Relativity. It allows you to build complex searches by rearranging and visualizing nested conditions by dragging and dropping them. Conditions can be easily set and organized into logic groups, with the logic display at the top of the panel updating automatically to reflect any changes made. You can access your dtSearch and keyword search indexes and fields within the workspace when creating a search.

With the search panel, the search conditions tool allows you to refine searches by selecting fields, operators, and values to customize queries. This tool can be used independently or in conjunction with keyword searches, dtSearches, Analytics, or Cluster visualization. See the Searching Guide for more information.

1.3 Keyword searches

You can run keyword searches from the Documents tab and from Dynamic Object tabs. With these searches, you can leverage the basic functionality for querying the SQL full-text index populated with data from extracted text fields. The keyword search engine supports the use of Boolean operators and wildcards. See the Searching Guide for more information.

1.4 Saved searches

Saved searches give you the functionality to define and store queries for repeated use. With flexible settings, you can create a saved search based on any Relativity search engine, assign security permissions to it, and define specific columns to display your search results. Saved searches support the development of complex queries that you build using a form with search condition options. These queries run dynamically to ensure that updated results appear when you access a saved search. See the Searching Guide for more information.

1.5 dtSearch

Available on the Documents tab, you can use dtSearch's advanced searching functionality to run queries with proximity, stemming, and fuzziness operators, along with basic features such as Boolean operators and wildcards.

System administrators can create a dtSearch index for a specific subset of documents in a workspace, and then assign security to it. They must manually update indexes when you modify the document search sets used to create them. See the Searching Guide for more information.

1.6 Analytics

Supporting conceptual searching, Analytics includes documents in a result set when they contain similar ideas or conceptual relationships, rather than matching specific search terms or conditions.

You can create searches with Analytics that categorize your documents based on the concepts contained in a sample document set. Instead of categorizing documents, you can also perform clustering, which uses specific algorithms, system-defined rules, to identify conceptually related documents. See the Searching Guide for more information.

1.7 Regular expressions


Regular Expressions (RegEx) is a form of advanced searching that looks for specific patterns, as opposed to certain terms and phrases. With RegEx you can use pattern matching to search for particular strings of characters rather than constructing multiple, literal search queries. You can use RegEx with a dtSearch index using dtSearch syntax options to construct complex queries. See the Searching Guide for more information.

1.8 Additional features



Relativity provides additional features that makes searching easy to use from the Documents tab. Use the **search bar** to run a keyword search query, or click the drop-down list to select another search index you created.

Keyword Search ▼	<input type="text" value="Enter Search Terms"/>	Search	Clear
------------------	---	------------------------	-----------------------

With the *search condition* option, you can build queries using the same condition options available for saved searches.

You can click **Save as Search**  on the Documents tab to create saved searches based on the criteria defined for keyword searches, dtSearches, Analytics, or the search conditions option. See the Searching Guide for more information.

You can also use the search bar to view recent searches. Click in the search bar to see up to 10 of the most recent searches from any index in descending order, with the most recently run search first.

Keyword Search ▼	<input type="text" value="Enter Search Terms"/>	Search	Clear
 There are no documents	1 search 3	<input type="text"/>	
	2 search 2		
	3 search 1		
	Clear Recent Searches		page 

Select any search from the list to run that search. To clear the list, click **Clear Recent Searches**.



2 Filters

Filtering provides a fast and easy way to search for items in a list in Relativity. You can use filters to search for values in the fields on the active view, and across all records available in the searchable set. Filters are also available for item lists on tabs and pop-up windows.

The field type associated with each column determines the available filter types, such as text boxes, pop-up pickers, and drop-down lists. You do not need any specific security permissions to use filtering.

2.1 Showing and hiding filters in the item list


You enable filters by clicking an icon located at the top of the document list. You can:

- **Show/Hide filters** ()—to display the field filters at the top of each column. This icon turns orange when you activate a filter. Click the icon again to hide the filters. Your filter settings remain unchanged.
- **Clear all** ()—to remove the current filter settings. This option is only available after you set a filter.


The filter type determines the steps required to set the filtering criteria that it uses. See [Filter types](#) for details on filter types and their use.

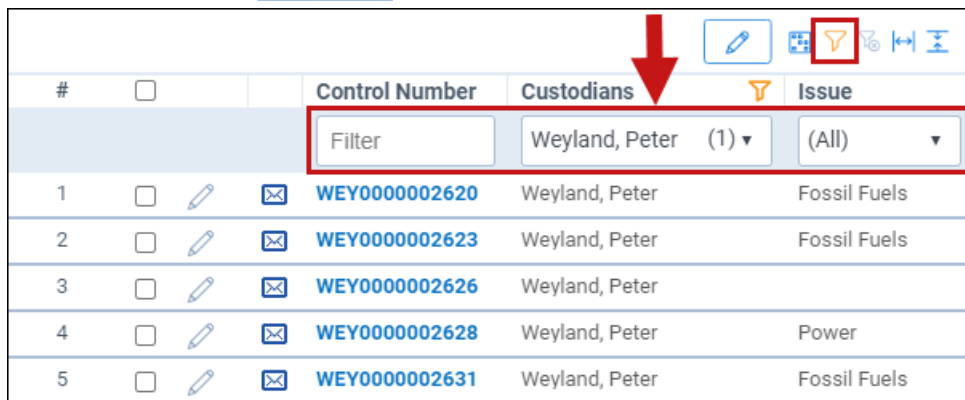
Note: When you create or edit a field, you can select its **Filter Type**. See Fields in the Admin guide.











After enabling filters, you are ready to enter criteria for the fields in your list.

Note: Before defining a new search filter, click the **Clear All** () icon or the **Clear All Conditions** link in the Search Conditions panel to remove any filter settings. For example, say you create a filter on the Privileged field, but now want to search only by the Custodian field. If you do not clear the filter, Relativity filters your documents by both the Privileged and Custodian fields, and your results may include fewer documents than you expect.

To use a filter:












1. If filtering is not enabled, do so by clicking the filter  icon to display the filtering fields below the column headings. The icon will change to orange.
2. Enter your filter criteria. The filter type determines how you enter the criteria and how to use operators to further define the filter. See [Filter types](#) for details on filter types and their use.



#		Control Number	Custodians	Issue
		Filter	Weyland, Peter (1) ▼	(All) ▼
1	<input type="checkbox"/>	  WEY0000002620	Weyland, Peter	Fossil Fuels
2	<input type="checkbox"/>	  WEY0000002623	Weyland, Peter	Fossil Fuels
3	<input type="checkbox"/>	  WEY0000002626	Weyland, Peter	
4	<input type="checkbox"/>	  WEY0000002628	Weyland, Peter	Power
5	<input type="checkbox"/>	  WEY0000002631	Weyland, Peter	Fossil Fuels

3. Press **Enter** to display the filter results.


The field column header displays the filter icon, letting you know which fields have filters applied. This persists whether you have the filters shown or hidden.

#	<input type="checkbox"/>		Control Number	Custodians	Primary Date/Time	Record Type 
			Filter	(All) ▼	(All) ▼	Email (1) ▼
1	<input type="checkbox"/>		 WEY0000002620	Weyland, Peter	02/16/2001 5:24 PM	Email
2	<input type="checkbox"/>		 WEY0000002621	Arnold, Johnny	02/16/2001 5:24 PM	Email
3	<input type="checkbox"/>		 WEY0000002622	Zipper, Andy	02/16/2001 5:24 PM	Email
4	<input type="checkbox"/>		 WEY0000002623	Weyland, Peter	02/16/2001 5:24 PM	Email
5	<input type="checkbox"/>		 WEY0000002624	Arnold, Johnny	02/16/2001 5:24 PM	Email

The Search Conditions panel also shows the filtering conditions.

Condition

1

List Conditions 

1. Record Type
any of these: Email

Clear All Conditions

Run Search

Auto-run ☐

You cannot edit the Search Conditions cards by clicking on them. However, you can clear all the conditions by clicking the link at the bottom of the panel.

2.2 Filter types

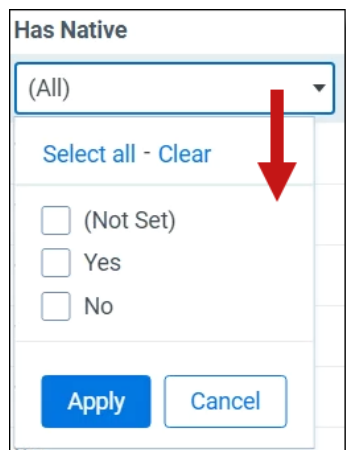
Relativity contains the following filter types:

- [Boolean](#)—Yes/No fields.
- [Numeric](#)—decimal, whole number, and currency fields.
- [List](#)—single choice, single object, user, and fixed-length text fields.
- [Date](#)—date fields.

- [Multi-list](#)—multiple choice, multiple object, and single choice fields.
- [Text box](#)—fixed-length text, long text, date, whole number, decimal, currency, and object fields.

2.2.1 Using Boolean filters

Boolean filters are available for Yes/No field types. To use a Boolean filter, click the drop-down menu and select **Yes**, **No**, or **(Not Set)** from the list. Click **Apply** to set the conditions. Click **Select all** to select all the options. Click **Clear** to remove the filter settings.



2.2.2 Using numeric filters for numbers

The numeric filter is available for the following field types:

- Decimal
- Whole Number
- Currency

To use a numeric filter, click the drop-down menu and select an operator. Operators include:

- Equal to (=)
- Not equal to (!=)
- Greater than (>)
- Less than (<)
- Less than or equal to (<=)
- Greater than or equal to (>=)

Enter a numerical value in the filter text box. After you enter the value in the text box, press **Enter** to apply the filter. For example, you might search for documents having a file size greater than 500000 bytes.

The screenshot shows a 'File Size' filter interface. At the top, there's a header 'File Size' with a funnel icon. Below it, a dropdown menu is open, showing a list of comparison operators: '>', '=', '!=', '<', '<=', '>', and '>='. The value '500000.00' is entered in the text box to the right of the dropdown. The '>' operator is currently selected and highlighted in blue.

When you click inside a filter text box, the **Advanced** link appears. Clicking this link opens the advanced filters screen for the filter type.

This screenshot shows the 'File Size' filter interface. The operator is set to '=', and the text box contains the word 'Filter'. Below the text box, there is a red-bordered button with a link icon and the text 'Advanced'.

Here, you can adjust the operator and add conditions to further define the filter. Click **Apply** to set the conditions.

The screenshot displays the 'Filter: File Size' configuration screen. It features two identical filter condition boxes stacked vertically, separated by the word 'OR'. Each box contains an 'Operator' dropdown set to 'Is' and a text input field with the placeholder 'Enter a number'. At the bottom left, there is a red-bordered button with a plus icon and the text '+ Add condition'. At the bottom right, there are two buttons: 'Apply' (in blue) and 'Cancel' (in white with a blue border).

2.2.3 Using List filters

The List filter is available for the following field types:

- Single Choice
- Single Object
- User
- Fixed Length Text

List filters are often associated with fields used for coding documents. The options displayed in the drop-down menu vary by the type and purpose of the associated field. For example, a field called Responsiveness might have the filter conditions of Responsive, Not Responsive, or Needs Further Review. Select the options that you want to filter for, and click **Apply**. Alternatively, you can click **Select all** to select all filters in the list. Click **Clear** to remove all selections. Click **(Not Set)** to show items where the field is empty (null).

The screenshot shows a user interface for filtering documents. At the top, there are two tabs: 'Responsiveness' and 'Issues'. Below each tab is a dropdown menu currently set to '(All)'. A red arrow points to the 'Select all - Clear' link in the dropdown menu for 'Responsiveness'. Below this link is a search bar labeled 'Search Items' with a magnifying glass icon. Below the search bar is a list of filter options, each with an unchecked checkbox: '(Not Set)', 'Non-Responsive', 'Responsive', 'Privileged', 'Needs Further Review', and 'Technical Issue'. At the bottom of the dropdown menu are two buttons: 'Apply' (highlighted in blue) and 'Cancel'.

If there is a single filter choice you want to apply immediately, hover over the item's row and click **Only** when it appears on the right side.

Record Type | Unified Title

(All) Filter

Select all - Clear ✕

Search Items 🔍

☐ (Not Set)

☒ Attach Only

☐ Edoc

☐ Email

Apply Cancel

Click the **Advanced filters** (🔗) icon to the right of the **Select all - Clear** options to launch the advanced filters dialog. This dialog has more operators and you can add or remove list items to the field's drop-down menu. Click **Apply** to set the conditions.

Filter: Responsiveness

Operator any of these

☒ Full Path

☐ Needs Translation

☐ Non-Responsive

☐ Password Protected

☐ Responsive

☐ Technical Issue/Corrupt File

Full Path

Filter

No data.

Apply Cancel

2.2.4 Using date filters

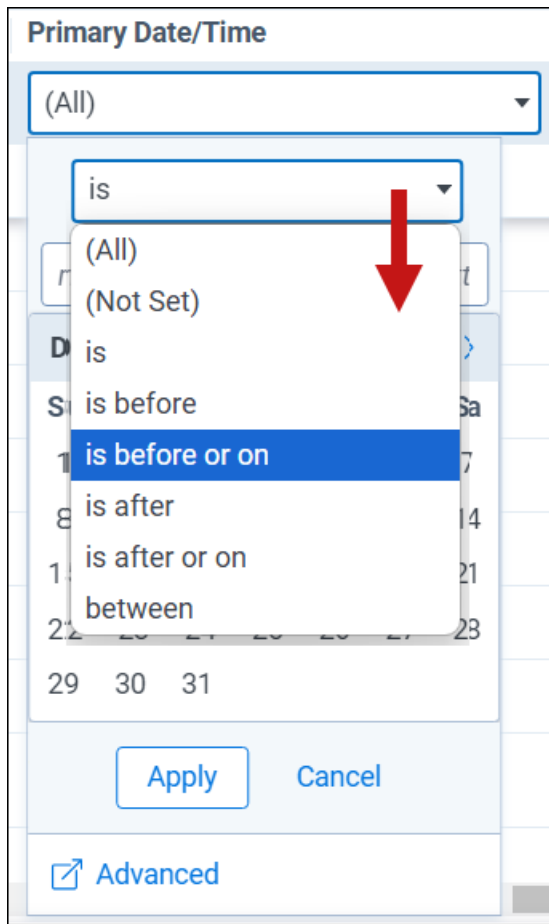
The date filter is available for Date field types to filter on dates and date ranges.

The following operators are available:

- (All)
- (Not Set)
- Is
- Is before
- Is before or on
- Is after
- Is after or on
- Between

For example, if you only want to view items sent before or on January 21, 2024 7 AM, do the following:

1. Select the **is before or on** operator.



2. Use the calendar interface to select the date, *January 21, 2024*. You can enter the date directly into the text boxes, or use the forward and back arrows to scroll through the calendar months.


Note: When you select the Between operator, two calendars display in the drop-down date filter, allowing you to pick the starting and ending dates and times.

3. Enter the time of 7:00 AM.

Note: Filtering on fields using the date and time format does not support using @Today. If you do not enter a specific time, Relativity enters the default time of 12:00 AM.

A calendar interface for filtering. At the top, a dropdown menu shows "is before or on". Below it are two input fields: "1/21/2024" and "7:00 AM". The calendar itself shows the month of January 2024, with the 21st highlighted in blue. At the bottom, there are "Apply" and "Cancel" buttons, and a link labeled "Advanced" with an external link icon.

4. Click **Apply** to set the conditions. Only items sent on or before 7 AM on January 21, 2024 appear in your document list.

Click the **Advanced** () link at the bottom of the calendar interface to launch the advanced filters screen. This screen has more operators where you can add additional date and time conditions. Click **Apply** to set the conditions.

The "Filter: Primary Date/Time" screen. It features a dropdown menu with "Is" selected. Below this are four input fields with placeholder text: "mm/dd/yyyy", "hh:mm tt", "mm/dd/yyyy", and "hh:mm tt". At the bottom, there is a "+ Add condition" link, and "Apply" and "Cancel" buttons.

2.2.5 Using multi-list filters

The multi-list filter is available for the following field types:

- Multiple Choice
- Multiple Object
- Single Choice

Using multi-list filters, you can select multiple conditions from a drop-down menu. These filters are often associated with fields used for coding documents.

To apply a multi-list filter, click the drop-down menu to display the conditions list. Select the conditions that you want to filter on. Click **Select all** to select all filters in the list. Click **Clear** to remove all selections. Select **(Not Set)** to show items where the field is empty (null). Click **Apply** to set the conditions.

If there is a single filter choice you want to apply immediately, hover over the item's row and click **Only** when it appears on the right side.

The screenshot shows a 'Record Type' filter dropdown menu. At the top, there's a '(All)' dropdown and a 'Filter' button. Below this, there's a 'Select all - Clear' link with a red arrow pointing to it. A search box labeled 'Search Items' is present. The list of filter options includes:

- ☐ (Not Set)
- ☒ Attach (highlighted row, with a red box around the 'Only' link to its right)
- ☐ Edoc
- ☐ Email

 At the bottom, there are 'Apply' and 'Cancel' buttons.

To find a specific list item in a large list, use the search box to enter a term. The list filters automatically as you type.


Record Type

Unified Title


(All) ▼

Filter

Select all - Clear



Search Items



☐ (Not Set)


☐ Attach

☐ Edoc

☐ Email

Apply

Cancel

Click the **Advanced filters** () icon to the right of the **Select all - Clear** options to launch the advanced filters screen. This screen has more operators where you can add or remove list items to the field's drop-down menu. Click **Apply** to set the conditions.

Filter: Record Type

Operator

any of these

any of these
none of these
is set
is not set

☐ Full Path
☐ Attach
☐ Edoc
☐ Email

>>
>
↔
<
<<

☐ Full Path
Filter
No data.

Apply
Cancel

2.2.6 Using text box filters

The text box filter is available for the following field types:

- Fixed-length Text
- Long Text
- Date
- Whole Number
- Decimal
- Currency
- Object

You can use text box filters to search on specific terms, numbers, and dates. Enter one or more terms in the filter text box. Connect multiple terms using the following operators:

- AND
- OR
- IS SET
- IS NOT SET
- BETWEEN

- = (equal)
- >= (greater than or equal to)
- <= (less than or equal to)

When entering terms and operators directly into a text box filter, you must use the proper syntax. See [Text box filter search examples](#) for more information.

You can also use the advanced text box filter feature to build more advanced filter queries on the selected text field. See [Using advanced text box filtering](#).

2.2.7 Text box filter search examples

Using operators, you can perform text searches by entering search strings directly into the filter text box. Additionally, you can use the advanced filtering feature to build more detailed queries. See [Using advanced text box filtering](#) for more information.

2.2.7.1 Search operators

The following table lists examples of valid search strings using search operators.

Valid search strings	Description of search results
cubs OR sox	Matches either <i>cubs</i> or <i>sox</i> or both.
cubs AND sox	Matches both <i>cubs</i> and <i>sox</i> . Both terms must appear for a successful match.
cubs OR sox AND kcure	Matches either <i>cubs</i> or <i>sox</i> or both, and <i>kcure</i> .
percent sign (%)	Acts similar to a wildcard, or the <i>is like</i> operator in a query. <i>%chard</i> matches <i>Richard</i> , <i>Pritchard</i> .
underscore (_)	Acts as a wildcard for a missing character. Do not use the underscore to check a field's value. It is slower and more resource-intensive than using the percent sign (%).
= with term	Matches an exact phrase.
cubs sox	Matches an exact phrase, <i>cubs(space)sox</i> .
IS SET	Returns only items where the field has a value.
IS NOT SET	Returns only items where the field does not have a value (null).

The following table shows examples of invalid search strings.

Invalid search strings	Description of search results
cubs AND	The AND operator requires a right search term.
cubs OR	The OR operator requires a right search term.
AND cubs	The AND operator requires a left search term.
OR cubs	The OR operator requires a left search term.

2.2.7.2 Alphabetical filtering

The following table lists examples of valid search strings you can use to filter text alphabetically.

Alphabetical filtering	Description of search results
>= c	Matches any term or phrase where the first letter is <i>c</i> or later in the alphabet.

Alphabetical filtering	Description of search results
<= c	Matches any term or phrase where the first letter is c or earlier in the alphabet.
= cubs sox	Matches the exact phrase <i>cubs(space)sox</i> .
cubs BETWEEN sox	Matches terms or phrases where the first letter of the phrase falls between c and s in the alphabet.

2.2.7.3 Dates and numbers

The following table lists examples of valid date and number searches, and the expected result set. When entering dates, do not place a zero (0) in front of single-digit values. For example, use 7/24/2024 and not 07/24/2024.

Valid search strings	Description of search results
>= 7/24/2024	Matches 7/24/2024 and later dates.
<= 7/24/2024	Matches 7/24/2024 and earlier dates.
= 7/24/2024	Matches the exact date of 7/24/2024.
>= 7/24/2024 1:23 PM	Matches 7/24/2024 1:23 PM and later dates and times.
<= 7/24/2024 1:23 PM	Matches 7/24/2024 1:23 PM and earlier dates and times.
= 7/24/2024 1:23 PM	Matches the exact date and time of 7/24/2024 1:23 PM.
7/24/2024 BETWEEN 8/24/2024	Matches dates that include and fall between 7/24/2024 and 8/24/2024.
7/24/2024 1:23 PM BETWEEN 8/24/2024 3:45 PM	Matches dates and times that include and fall between 7/24/2024 1:23 PM and 8/24/2024 3:45 PM.
7/24/2024	Matches the exact date of 7/24/2024.
>= 100	Matches numbers that are greater than or equal to 100.
<= 100	Matches numbers that are less than or equal to 100.
= 100	Matches 100 exactly.

The following table includes examples of invalid data and number search strings.

Invalid search strings	Description of search results
> 7/24/2024	You must use the equal sign with the greater than operator (>=).
< 7/24/2024	You must use the equal sign with the less than operator (<=).
>= 0/24/2024	The search string starts with 0 for the month. Instead of 01/24/2024, use 1/24/2024.
= 0/24/2024	The search string starts with 0 for the month. Instead of 01/24/2024, use 1/24/2024.
07/24/2024 BETWEEN 8/24/2024	The search string starts with 0 for the month in the starting date. Instead of 01/24/2024, use 1/24/2024.
7/24/2024 BETWEEN 08/24/2024	The search string starts with 0 for the month in the ending date. Instead of 01/24/2024, use 1/24/2024.

2.2.8 Using advanced text box filtering

You can use the following operators with advanced filtering:

- Is
- Is not
- Is set
- Is not set
- Is less than
- Is greater than
- Is like
- Is not like
- Contains
- Does not contain

To use advanced text box filtering:

1. Click inside the filter text box for a column. The **Advanced** link appears.
2. Click the **Advanced** link.
3. Choose an operator (such as is greater than.) See Fixed-length, long, or extracted text operators in the Searching guide for a list of definitions of the available operators.

The screenshot shows a dialog box titled "Filter: Email From". Inside, there is a light blue shaded area containing a dropdown menu with "is" selected and a large text input box with the placeholder "Enter text query". Below the shaded area, there is a blue link "+ Add condition", a blue "Apply" button, and a white "Cancel" button with a blue border.

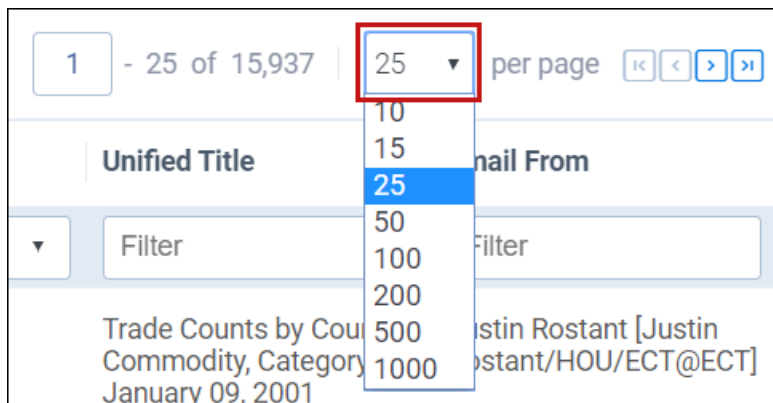
4. Enter a value into the query text box(such as 100.) The filter returns items where the value of the field is greater than 100.
5. (Optional) Click **Add condition** to add a new filter condition. For example, you may want to also return items where the text field contains the word privilege.

Note: Multiple conditions are automatically connected with an OR operator.

6. Click **Apply** to set the conditions.

2.3 Changing item sets per page

You can use the set selector menu to change the number of items that appear per page. The set selector menu is at the top of the document list. The option you select remains the default setting during your session until you select another option.



2.4 Saving filters as a search

To save your filtered item set as a saved search:

1. Click the **Save Search** button located at the bottom of the screen.
2. Select or enter the following required information:
 - **Name**—enter a title for the search. The title appears in the saved searches browser.
 - **Owner**—select **Public** to make the search available to all users or choose a specific user from the list. Click **Me** to select your name from the list, making the search private. (You must have the appropriate privileges to view searches.) See Controlling the visibility of saved searches in the Searching guide.
 - **Search Folder**—click the **Select** button to launch a pop-up window where you can save the new search to a specific folder on the saved searches browser. Highlight the folder where you want to save the search, and click **OK** to select it.
3. Add to or change the search criteria as needed. See Create and edit a saved search.
4. Click **Save**.

3 Keyword search

You can use a keyword search to query a full text index. The long text and fixed-length text fields included in this index vary by workspace.

Note: New workspaces created in RelativityOne have extracted text automatically stored in Data Grid. Workspaces restored into RelativityOne using the ARM application will automatically have the extracted text migrated to Data Grid. In order to search extracted text in workspaces, you must use dtSearch or Analytics searching; you cannot use keyword search.

You can use the AND, OR, NOT Boolean operators in keyword searches, as well as quotation marks for exact matches, asterisks (*) for wildcards, and other features. However, if you perform a keyword search with multiple terms, documents where those terms exist in separate fields won't return.

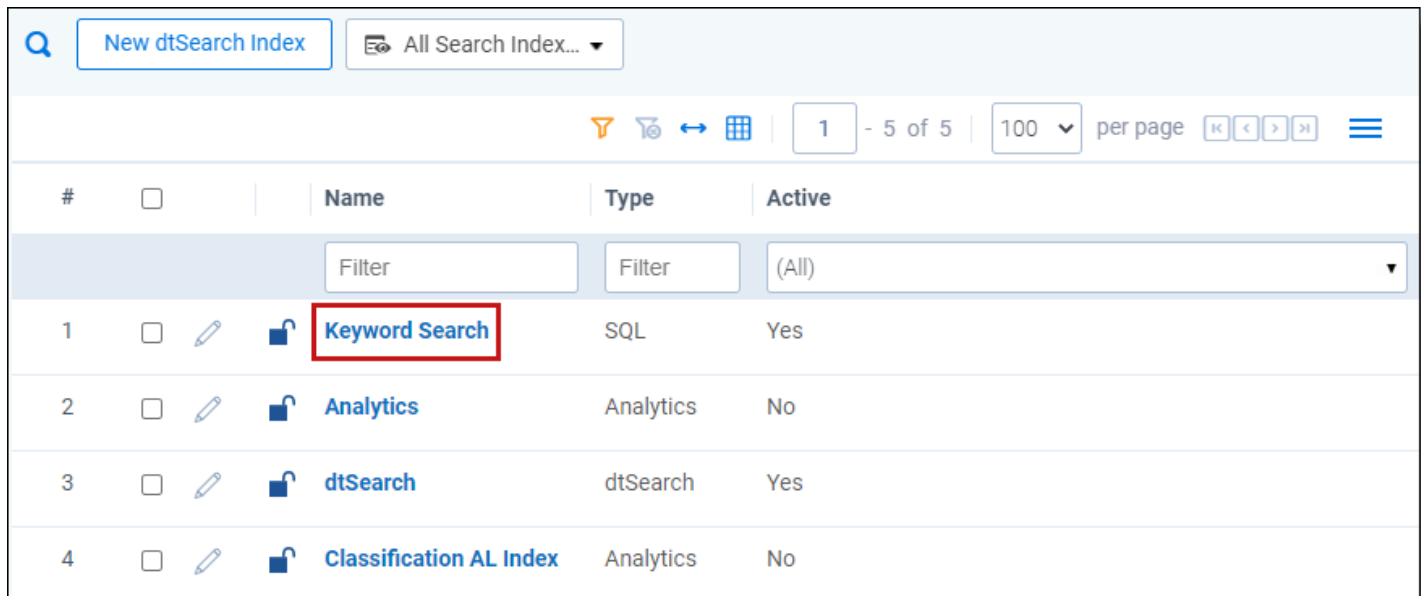
Note: If you want to draft queries outside of Relativity, use a plain text editor such as Microsoft Notepad to prevent adding characters or formatting that might return unexpected search results.

While the keyword search offers fewer options than other Relativity searches, it uses an index that's automatically populated, reducing maintenance and ensuring all required document fields are indexed.

Note: For information on configuring and managing word breakers, see [Microsoft's documentation](#).

3.1 Fields

A keyword search index is available in the Search Indexes tab by default. Click the **Keyword Search** link.



<input type="text" value="New dtSearch Index"/>				<input type="text" value="All Search Index..."/>			
				1 - 5 of 5 100 per page			
#	<input type="checkbox"/>	Name	Type	Active			
		<input type="text" value="Filter"/>	<input type="text" value="Filter"/>	<input type="text" value="(All)"/>			
1	<input type="checkbox"/>	Keyword Search	SQL	Yes			
2	<input type="checkbox"/>	Analytics	Analytics	No			
3	<input type="checkbox"/>	dtSearch	dtSearch	Yes			
4	<input type="checkbox"/>	Classification AL Index	Analytics	No			

The keyword search index details page contains the following fields:

- **Name**—the name of the keyword search index. The name is the display name for the index.
- **Order**—a number that represents the position of the index in the list. The lowest-numbered index will be at the top. The highest-numbered index is at the bottom. Items that share the same value are sorted in alphanumeric order. Index order can be any integer, positive or negative. No decimals are allowed.

- **Active**—determines whether the index should be activated or deactivated. **Yes** means that the index will be activated; **No** means that the index will be deactivated.

Note: If you apply item-level security to a search index, users cannot run any public saved searches built on that index and will get an error. We recommend leaving the index unsecured and instead applying security to the Search indexes tab or to individual saved searches.

3.2 Example keyword search strings

The following table lists search string examples with their expected results.

Search string	Returns documents with...
wired	the word wired
wired magazine	the words wired and magazine
wired AND magazine	the words wired and magazine
wired OR magazine	the word wired or the word magazine
wired, magazine	the word wired or the word magazine
"wired magazine"	the exact phrase wired magazine
wired NOT magazine	<p>the word wired and not the word magazine</p> <hr/> <p>Note: Do not start key word searches with the NOT operator, or use it with the OR operator. For example, these searches are invalid:</p> <ul style="list-style-type: none"> • not wired • wired or not magazine <p>See NOT Operator Evaluation in Keyword searches.</p>
wire*	<p>any words beginning with wire, such as wired, wires, wireless</p> <hr/> <p>Note: Key word searches do not support the use of wildcards at the beginning of a word. Keyword searches are SQL index searches run on the Microsoft SQL Server, which does not support leading wildcards in full text searches.</p>
computer AND (wired OR magazine)	<p>the word computer and the word wired OR the word computer and the word magazine</p> <hr/> <p>Note: When a search string does not include parentheses, the order of precedence for a keyword search evaluates AND then OR expressions. For example, the search string A AND B OR C is evaluated as (A AND B) OR C.</p>

Note: Search terms with accented letters are recognized and return keyword search results.

3.3 Using the NOT operator in keyword searches

When running a keyword search that is an SQL full text search, carefully format queries that use the NOT operator. For example, you may want to query for email messages that have Ryan as the author, but do not have Will as the recipient. The fields in the following record are included in the index used to demonstrate how this query is run:

Document	OCR	Recipient	Author
AS00001	From: Ryan To: Will	Will	Ryan

A keyword search using the string Ryan NOT Will returns the document AS00001 even though you would not expect it in the result set. The following table illustrates the SQL logic used to evaluate the query Ryan NOT Will.

SQL queries this field...	Returns these results...
OCR Field	Finds both Ryan and Will, so no document is returned.
Recipient Field	Does not find Ryan, so no document is returned.
Author Field	Finds Ryan but not Will, so the document AS00001 is returned.

When these fields are searched using the SQL logic, the Author field matches the query Ryan NOT Will, and unexpectedly returns the document.

Note: You can use the AND NOT operator in a dtSearch as an alternative approach to this type of keyword search.

3.4 Understanding noise words

Noise words in a keyword search include punctuation marks, single letters, single digits, and words such as "at", "a", "on" and "the".

3.4.1 Keyword search noise words - without double quotes

Noise words used in keyword searches are ignored if the search string is **not surrounded by double quotes**. In a search for the phrase *sun on my head*, both *on* and *my* are ignored. The result is that the words *sun* AND *head* are queried without respect to proximity. Thus, any documents that contain both the words *sun* and *head* will be returned.

The following table illustrates how keyword search queries for phrases that contain noise words that are not surrounded by double quotes.

Searching string without quotes	Queries for this string
sun on my head	sun AND head
sun on head	sun AND head

3.4.2 Keyword search noise words - with double quotes

If a keyword search string containing noise words is **surrounded by double quotes**, then the noise words' positions in the string are taken into account when the query is executed. However, only the *positions* of any *intervening* noise words are taken into account, not the noise words themselves. Noise words at the beginning or tail end of a keyword search string are ignored.

For example, the search strings "sun on my head" and "sun my on head" (where *on* and *my* are switched) return the same records. This is because keyword search evaluates both search strings as a query for the phrase **sun ABC XYZ head**, where ABC and XYZ represent **any two words**, not just noise words. Similarly, a query for the search string "sun on head" returns documents that contain the phrase **sun ABC head**, where ABC represents any word.

The following table illustrates how keyword search queries for phrases that contain noise words, and that are surrounded by double quotes.

Searching string (with quotes)	Queries for this string
"sun on my head"	sun [AnyWord] [AnyWord] head
"sun on head"	sun [AnyWord] head
"sun on my head and"	sun [AnyWord] [AnyWord] head
"and sun on head"	sun [AnyWord] head

3.4.3 Single digits as noise words

Single digits 0-9 are default noise words, so you cannot query on them with a keyword search. Relativity does not return the expected results if you attempt to query on a single digit. Use the `dtSearch` feature to query on a specific number or letter.

However, you can use a keyword search to query on whole numbers greater than 9. You can search on more than one digit, such as 09. While these digits may be used to represent a specific numeric value, such as 9, they are not considered single digits, and can be used in a keyword search.

3.4.4 Punctuation as noise words

Certain punctuation marks are treated as noise words by default, so you cannot query on them with a keyword search. They include:

- Period (.)
- Colon (:))
- Semicolon (;)
- Slash (\,/)

3.4.5 At sign (@) and dashes

The at sign (@) and dashes (-) are ignored from being indexed in a keyword search, when either is used at the beginning of a query. For example, if you search a domain name, the same number of documents return whether you include or exclude @.

3.4.6 Hyphens and dashes

When a search phrase includes a hyphen or dash, the query returns results that include terms containing other punctuation marks. For example, the following results return for a search on the term Pop-up:

- Pop.up
- Pop--up

3.4.7 Default noise word list

Relativity comes with the following default noise words:

Begins with...	Noise words
A	about, after, all, also, another, any, are, as, at
B	be, because, been, before, being, between, but, both, by
C	came, can, come, could

Begins with...	Noise words
D	did, do, does
E	each, else
F	for, from
G	get, got
H	has, had, he, have, her, here, him, himself, his, how
I	if, in, into, is, it, its
J	just
L	like
M	make, many, me, might, more, most, much, must, my
N	never, no, now
O	of, on, only, other, our, out
S	said, same, see, should, since, so, some, still, such
T	take, than, that, the, their, them, then, there, these, they, this, those, through, to, too
U	under, up, use
V	very
W	want, was, way, we, well, were, what, when, where, which, while, who, will, with, would
Y	you, your

3.5 Running a keyword search


3.5.1 Running a keyword search in the search panel

Use the following steps to run a keyword search in the search panel.

1. Navigate to the Search panel in the Documents Tab.
2. Click **Add Condition**.
3. Select **(Index Search)** in the Add Condition drop-down menu. The (Index Search) window opens.
4. Select **Keyword Search** from the drop-down Index.
5. Enter terms for the search in the **Search Terms** box.
6. Optionally, select the **Sort By Rank** option to return results in order by relevance. The most relevant documents are listed at the top of the result set.
7. Click **Apply**.
8. (Optional) Add any additional conditions through the Add Condition drop-down menu.
9. Click **Run Search**. To stop a long running search, click **Cancel**.

3.5.2 Running a keyword search in the Search browser

Use the following steps to run a keyword search in the Search browser.

1. Click  to access the search browser from the document list.
2. Click **New Search**.
3. Set required fields.
4. Click **Add Condition**.
5. Select **(Index Search)** in the Add Condition drop-down menu. The (Index Search) window opens.
6. Select **Keyword Search** from the Index drop-down menu.
7. Enter terms for the search in the **Search Terms** box.
8. Optionally, select the **Sort By Rank** option to return results in order by relevance. The most relevant documents are listed at the top of the result set.
9. Click **Apply**.
10. (Optional) Add any additional conditions through the Add Condition drop-down menu.
11. Click **Save** or **Save As**.
12. Click the name of the keyword search in the search browser.
13. Click **Run Search**. To stop a long running search, click **Cancel**.

4 Search panel

The search panel is available for the Document list and for many other tabs in Relativity.

Using the search panel, you can build complex searches using drag-and-drop to rearrange and visualize nested conditions. You can easily set conditions and drag and drop them into logic groups. The logic display at the top of the panel updates automatically to reflect your drag-and-drop changes. You can access your dtSearch and keyword search indexes and fields in the workspace when creating your search.

4.1 Permissions

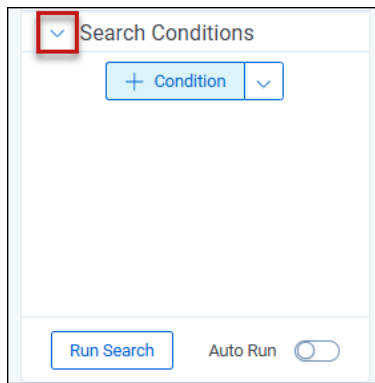
Search panel functionality is controlled by security permissions. If neither permission below is present, the search panel buttons will not appear.

- To access search indexes (Keyword search, dtSearch, Analytics) from the Add Conditions drop-down menu, you must have the View permission enabled for Search Index.
- To view and select fields from the Add Conditions drop-down menu, you must have the Edit permission enabled for Search.

See Workspace security for more information on setting security permissions.

4.2 Collapsing, expanding, and moving the search panel

You can collapse the search panel by clicking the down arrow and expand it again by clicking the right arrow (>) next to the collapse panel.




To move the search panel to either side of the document list, click on and hold the top of the panel until the cross



arrows appear () and drag it to the new location.

4.3 Condition card icons

The condition card features the following icons and menu options:



Name	Icon	Description
Grabber	...	Click this icon to drag the condition card to a new location.
Expand/Collapse	▼	Click this icon to expose all of the filters and conditions for the card. Collapse the card to return to the summary display.
Clear	×	Click this icon to clear all of the card's filters and conditions.
Copy	📋	Click this icon to copy the card's filters and conditions. You can paste the card within the pane, then make your adjustments.
Enable/Disable	<input checked="" type="checkbox"/>	Click this box to enable the card's filters and conditions. Uncheck this box to disable the card's filters and conditions without actually deleting the card itself.

4.4 Using the auto-run search setting


Enable the auto-run setting at the bottom of the searching panel to automatically update your data when you select new conditions.

With Auto-run search set to **On**, each condition that you select from the searching panel will cause the page to auto-refresh and display updated data based on the new condition. You can leave Auto-run Search set to **Off** if you prefer to manually apply any new conditions as you add them to your search.

The screenshot shows a search panel interface. At the top, there is a '+ Condition' button and a dropdown arrow. Below this, the number '1' indicates one condition is active. A condition card is displayed with the text '1. Has Images' and 'any of these: Yes'. Below the card is a 'Clear All Conditions' link. At the bottom of the panel, there is a 'Run Search' button and an 'Auto-run' toggle switch, which is currently turned on and highlighted with a red box.

4.5 Creating a search in the search panel

To create a search using the search panel:

1. Expand the search panel from the item list by clicking the  icon in the upper left corner of your screen.
2. Click **Add Condition** in the search panel.

Note: If a related items condition has been set for the saved search, the related items condition also applies to any conditions that are added to the saved search via the search panel.

3. Do one of the following:
 - Enter or select the field to which you want to apply the condition from the drop-down menu.
 - If you have previously linked fields to a field category, you can alternatively select the desired field category from the drop-down menu and then enter or select the field to which you want to apply the condition from the list. To learn more, visit [Field Categories](#).
4. Depending on the field, you will be prompted to specify the conditions to add.

Note: For single object and multiple object fields, the name of the objects acts as hyperlinks in the filter condition card. You can quickly navigate to the named object instance by clicking on it. For example, if you have a search condition that names the "First pass review" batch set, you can navigate directly to that batch set from the filter condition card.

5. (Optional) Add additional conditions.
6. (Optional) Apply logic groups to your conditions. See [Applying logic groups](#).
7. Click **Run Search** if auto-run search is toggled off). To cancel a long running search, click **Cancel**.

The search will be automatically saved and you can reference it later as a saved search.

Note: To remove all conditions from the search panel click **Clear All Conditions**.

If you want to edit a condition, click on the search card.

Note: When you search on a user-created date field using a relative date, Last 7 Days for example, you may see different results between a saved search and a search from the search panel if you are not in the same timezone as your Relativity instance.

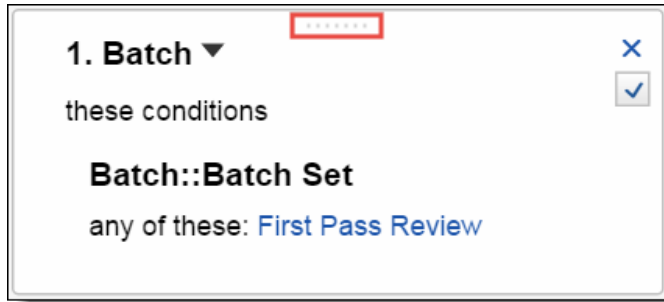
4.6 Applying logic groups to search conditions

Logic groups act as visual parentheses for your search query. The criteria within logic groups are evaluated first before evaluating against other search conditions or logic groups. When creating logic groups, the logic display at the top of the search panel updates automatically to reflect your drag-and-drop changes.

To apply logic groups to search conditions:

1. Click **Add Logic Group**.
A green frame appears.

2. Click the handle on the top of the filter condition card you want to add to a logic group.



3. Drag the condition into the logic group frame.
4. Add other conditions to the logic group as needed.

Note: You can also create a logic group automatically by dragging one condition onto another.

5. Click the **AND** or **OR** drop-down menus to set your operators inside your logic group.
6. (Optional) Add additional logic groups and repeat steps 3-5 for the logic groups you add.

7. Click **Run Search** if auto-run search is toggled off. To cancel a long running search, click **Cancel**.


The screenshot shows a search panel interface. At the top, there is a button labeled "+ Condition" with a dropdown arrow. Below this, the main condition is displayed as "(1 AND 2) OR (3 AND 4)". The panel is organized into two main sections, each enclosed in a light green border. The first section contains two conditions: "1. Batch" with sub-conditions "these conditions" and "Batch::Batch Set any of these: First Pass Review", and "2. Has Images" with sub-conditions "any of these: Error; No". These two are connected by an "AND" operator. The second section contains two conditions: "3. Batch" with sub-conditions "these conditions" and "Batch::Batch Set any of these: First Pass Review", and "4. Has Images" with sub-conditions "any of these: Yes". These two are also connected by an "AND" operator. The two main sections are connected by an "OR" operator. At the bottom of the panel, there is a button labeled "Clear All Conditions".

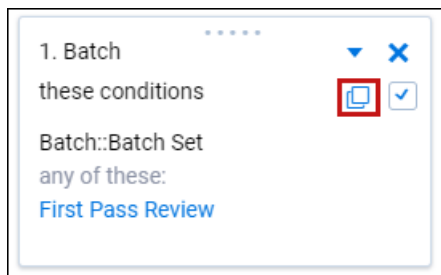
Note: To remove all conditions from the search panel, click **Clear All Conditions**.

4.7 Navigating the search panel

Within the search panel you can expand and collapse your search conditions, in addition to toggling the conditions on and off. By using these options you can better visualize data and make quick changes.

4.7.0.1 Copying conditions

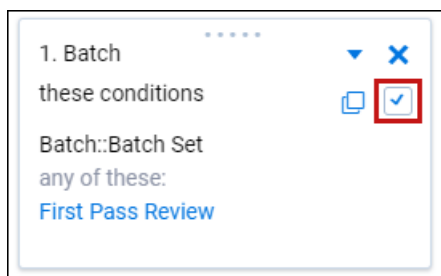
To copy a condition, click the Copy  icon in the search panel card.



A pop-up modal opens. In the modal, click **Apply** to copy the same search condition. Edit the operator and the search before clicking Apply for a variation of the original search card. Click **Add Condition** to add another condition to your new search card.

4.7.0.2 Toggling conditions on and off


You can toggle conditions within the search panel to change the documents the search returns. To toggle conditions on and off click the checkbox in the lower-right corner of the condition box. If the box is checked the condition is on; if the box is unchecked the condition is off.

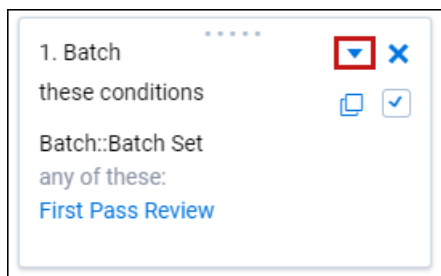


Click **Run Search** to update the toggled conditions (if auto-run search is toggled off).

Note: To collapse and expand filter condition cards in the panel click the arrow next to the name of the field applied in the condition.

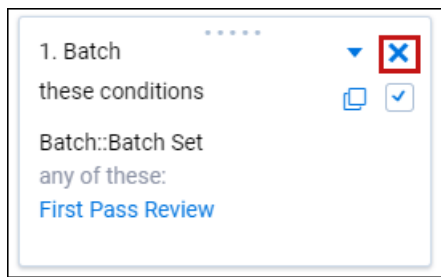
4.7.0.3 Expanding/collapsing cards

If you have a very long search panel card, you can collapse the card. To collapse a card, click  in the top right corner of the search panel card.



4.7.0.4 Removing conditions

To remove a condition, click **X** in the upper right corner of the condition box.



Click **Run Search** to update the toggled conditions (if auto-run search is toggled off).

5 Searching workflows

Relativity provides flexibility to help you identify documents through searches so that you can ready them for further review and analysis.

This topic highlights setting up multiple search indexes, applying various search techniques, and using filters properly when executing searches.

Note: For details on searching for dates within Relativity using a filter, a saved search, or a dtSearch, see the [Searching for dates in Relativity](#) knowledge base article on the Relativity Community.

5.1 Workflow details

Relativity offers both Basic Keyword Search and Specialized dtSearch.

The Basic Keyword Search Index populates automatically, is available out of the box, and supports basic Boolean operations. These operations include AND, OR, NOT, and the wildcard (*) operator.

The Specialized dtSearch Index is custom built and must be set up with an index to query against. In addition to basic Boolean operations, it also supports the following functionality:

- Wildcard operator for a single digit (=)
- Stemming operations (~)
- Fuzzy search operation (%)
- Customize data set to be indexed via custom saved search
- Proximity Searching
- Customized Noise Word List
- Auto Recognition of email addresses, credit card numbers, and dates.

5.2 Best practices for advanced operators

The following includes best practices for the use of advanced operators and workflow options:

5.2.1 Proximity search

Proximity search uses operators to search certain terms in proximity to other terms in a document. Use the following recommendations when employing proximity searching:

- To use directional proximity searching use "pre /x " operators. To use non-directional proximity searching use "w /x" operators.
- To determine the beginning and end of a document, use reserved word with either ("xfirstword") or ("xlastword"). Use these operators to search for metadata like email addresses or footers within a document.
- To determine proximity. The distance between terms is important so you know whether you receive true hits or

false positives. The following image illustrates the proper syntax for proximity searches:

Searching for → **Blue Orange Yellow Red Green Brown**

Distance → 1 2 3 4 5

Blue w/1 Orange w/1 Yellow w/1 Red w/1 Green w/1 Brown

INCORRECT

((((Blue w/1 Orange) w/2 Yellow) w/3 Red) w/4 Green) w/5 Brown)

CORRECT

The following image shows all correct and incorrect combinations of proximity searches that result in successful hits when using proximity operators:

✓	Blue		Red	✓	Blue		"Little Red Riding Hood"
✓	(Blue or Orange)		Red	✓	(Blue or Orange)		"Little Red Riding Hood"
✓	(Blue w/n Orange)		Red	✓	(Blue and Orange)		"Little Red Riding Hood"
✓	((Blue or Orange) w/n Yellow)		Red	✓	(Blue w/n Orange)	W/n or PRE/n	"Little Red Riding Hood"
✓	((Blue or Orange) w/n Yellow)		(Red or Green)	✓	((Blue w/n Orange)		Red) w/n Green)
✓	(Blue and Orange)	W/n or PRE/n	Red	✗	(Blue w/n Orange)		(Red w/n Green)
✓	(Blue and Orange)		(Red or Green)	✗	(Blue and Orange)		(Red and Green)
✓	Blue%%		Red%%	✗	(Blue w/n Orange)		(Red and Green)
✓	(Blue and Orange)		Red%%				
✓	Blue~		Red~				
✓	(Blue and Orange)		Red~				

Sometimes additional factors affect the distance between terms. The following image demonstrates a search for a string of words in proximity to another word. The search in the following image takes the following into account:

- Noise words count as words when calculating proximity.
- Punctuation counts as whitespace when using default settings.
- Relativity treats line breaks and consecutive space characters as single spaces.
- dtSearch default noise words and connector words like AND, OR, BETWEEN, and NOT count as words when calculating proximity.

The **message**, together with any attachment, may contain confidential and/or privileged information.

In this case, we continue to calculate the distance of the string ("confidential and/or privileged information:") from the word ("message") as shown above. However, when using connector words, system admins should create an index that removes noise words from the noise word list.

Noise Words
a
about
after
all
also
an
and
another
any
are

Relativity reserves the following noise words and characters, which continue to behave as operators, as well as being noise words: and, or, not, to, contains, xfirstword, xlastword, ", (), *, ?, %, @, ~, #, &, :, =.

Once that index is available for query you can either place the search string in quotes or apply stemming to the connector words to override their function as a connector.

**(i) message pre/7 "confidential and/or
privileged information"**

OR

**(ii) message pre/7 confidential and~/or~
privileged information**

5.2.2 Auto-Recognition

Auto-Recognition identifies email addresses, dates, and credit card numbers in the data set you want to index. You can turn this feature on and off when you build your index. The search returns the results regardless of the data format. Auto-Recognition adds some time to your index build, but, depending on your case, the benefits can be significant.

Type	Syntax	Results	Benefits
Email	mail(*doe@kcura.com)	jdoe@kcura.com john.doe@kcura.com	Treats email addresses as a single term regardless of alphabet settings for "@" and "." characters. Can be used with wildcards.
Credit Card Numbers	Creditcard(123467813246578)	1234-5678-1234-5678 1234 5678 1234 5678 1234567812345678	Returns sets of numbers that satisfy the criteria for valid credit card numbers. Can be used with wildcards.
Date	date(1/15/2006)	January 15, 2006; 15 Jan 06; 2006/01/15; 1/15/06; 1-15-06; The fifteenth of January, two thousand six	Returns date regardless of format. Dates mentioned in body of documents and not in metadata.

5.2.3 Searching for times

If you want to search for times in the body of documents, perform a full-text search with your dtSearch index for a specific time. Keep in mind that some characters cause a word break, such as the colon and period. Searching for 12:15 p.m. results in searching for four words: 12, 15, p, and m.

5.2.4 Filters

In addition to the Search Indexes, you can also use filters to search on metadata fields and narrow down the review set of documents.



The following list includes metadata fields, their corresponding field types, and the filter type available in Relativity:

- **File Type** - Single-Choice, List Filter
- **Custodian** - Single-Choice, Pop-Up Filter
- **Date Sent** - Date, Textbox Filter
- **Email To** - Long Text, Textbox Filter
- **Email Subject** - Fixed Length Text, Custom Filter
- **Designation** - Single Choice, Multi-Choice List

The following table shows different types of filters you can set up with available fields in Relativity.





























































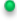

















Type	Text Box	Single Choice List	Multi-Choice List	Pop-Up	Custom
Fixed Length Text	✓	✓			✓
Long Text	✓				
Date	✓				✓
Number (Whole, Decimal, Currency)	✓	✓			✓
Yes/No		✓			
Choice (Single, Multiple)		✓	✓	✓	
Object (Single, Multiple)	✓			✓	

5.2.5 Troubleshooting workflow

To troubleshoot and test searches, use a white board approach to map out searches before you run them. As part of the process, try to take into account all possible variations, so you can see which documents Relativity returns and to gain a better understanding of the search. For instance, if you want to find email family groups with inconsistent coding, numerous possibilities for searches exist. The following illustration shows a mapping of nine email family groups, each with an email and two attachments.

Construct searches (1 and 2) and mark the corresponding documents, and their family members, with hits for each of the searches. At this point, when you look at the board, you can see that the four email family groups have one thing in common. The documents are responsive to both searches. So, to close it out, construct a third search that pulls back documents common to both searches.

Conflict Check within Family Groups

 	DOC000001	Responsive	 	DOC000010	Privileged	   	DOC000019	Privileged
 	DOC000002	Responsive	 	DOC000011	Privileged	   	DOC000020	Responsive
 	DOC000003	Responsive	 	DOC000012	Privileged	   	DOC000021	Responsive
 	DOC000004	Not Responsive	   	DOC000013	Responsive	 	DOC000022	
 	DOC000005	Not Responsive	   	DOC000014	Not Responsive	 	DOC000023	
 	DOC000006	Not Responsive	   	DOC000015	Not Responsive	 	DOC000024	Not Responsive
 	DOC000007		   	DOC000016	Responsive	   	DOC000025	
 	DOC000008		   	DOC000017	Privileged	   	DOC000026	Responsive
 	DOC000009		   	DOC000018	Responsive	   	DOC000027	

1. Designation is Responsive + Include Family
2. Designation is Not Responsive OR Privileged OR Not Set + Include Family
3. 1 AND 2

6 Search conditions

You can use the search conditions option to build complex queries by selecting fields, operators, and values. While this feature has the same functionality as the search condition section of the saved search form, it's conveniently available from the Documents tab and Relativity Dynamic Object tabs. This option displays up to five rows, with each row representing a separate criterion. Depending on the type of field you select, different operators appear. You can use this option alone or in conjunction with keyword searches, dtSearches, Analytics, or Cluster visualization. When you use search options in conjunction with another search feature, documents must both meet the search criteria and also the conditions specified.

Note: To use the search conditions option, you must have add or edit permissions for **Search** and access to the **Saved Searches Browser** assigned to you through the Security page. See Workspace security in the Admin Guide.

6.1 Setting up search conditions

6.1.1 Setting up search conditions in the Search panel


To set up search conditions in the Search panel follow these steps:

1. Navigate to the Search Builder.
2. Click **Add Condition** and select the fields to which you want to apply conditions.
A pop-up window opens for each condition field you select.
3. Set the required conditions in the relevant field pop-up window by selecting the desired field or by choosing a field category and then selecting the desired field.
4. Click **Apply**.
5. (Optional) Click **Add Logic Group** to add a logic group. Logic groups are evaluated first, and then connected to other filter conditions or logic groups using AND / OR operators.
6. (Optional) Drag and drop conditions together to create logic groups.
7. (Optional) Add the AND or OR operators to connect the criterion.
8. Click **Apply**.

If you need to edit the condition, click on the condition card. The pop-up reopens so you can make changes.

6.1.2 Setting up search conditions in the Search browser

To set up search conditions in the Search browser, follow these steps:

1. Click the **Saved Search** icon to navigate to the Search Browser.

2. Click **New Search**.
3. Click **Add Condition** and select the fields to which you want to apply conditions.
A pop-up window opens for each condition field you select.
4. Set the required conditions in the relevant field pop-up window.
5. Click **Apply**.

6. (Optional) Click **Add Logic Group** to add a logic group and drag and drop your conditions into the frames. Logic groups are evaluated first and then connected to other filter conditions or logic groups using AND / OR operators.
7. (Optional) Add the **AND** or **OR** operators to connect the criterion.
8. Enter all required fields.
9. Click **Save** or **Save As**.

If you need to edit the condition, click on the condition card. The pop-up reopens so you can make changes.

6.1.3 Using the multiple object condition builder

You can set conditions for **multiple object fields** from a single pop-up window. The following example uses Batch as a multiple object field condition.

To use the multi-object condition builder, follow these steps:

1. Click **Add Condition** from the Search browser or Search panel.
2. Enter or select **Batch**.
The Condition: Batch window opens.
3. Select the desired operator from the **Operator** drop-down list. For example, **these conditions**.
4. Click **Add Condition**.
5. Select the condition from the **Add Condition** drop-down list. For example, **Batch:: Batch Set**.
6. Click the available fields from the **Available** column and move them to the **Selected** column as desired.
7. Click **Apply**.
8. (Optional) Click **Add Logic Group** to add a logic group. Logic groups are evaluated first and then connected to other filter conditions or logic groups using AND / OR operators.
9. (Optional) Drag and drop conditions into logic groups as desired.
10. Click **Apply** on the Condition: Batch window.
Relativity applies your search conditions.

For more information about multiple object searching logic, see Multiple object searching.

6.2 Canceling queries

You can cancel a long-running search or view by clicking the **Cancel** button. This link appears when you perform a keyword search, dtSearch, Analytics search, or use Pivot. It also appears when you filter or sort a document or Dynamic Object list. It also appears when you perform other actions on item lists containing documents or Dynamic Objects that initiate a query in the background.

When you click **Cancel Request**, Relativity stops the background query used to populate documents in an item list. If you edit a search and click this link, your changes save, but the item list doesn't load. For example, your changes save when you perform a mass edit on a list of documents in a search and then click **Cancel Request** when the query is running to redisplay the updated list.

Note: Relativity creates an audit record in the History tab for canceled queries. The query description displays the running time of the query and indicates that it was canceled. You must have the appropriate permissions to view this tab.

The following table explains different scenarios in which you might cancel a running query and whether the query actually cancels.

Scenario	Result
You start a query and click Cancel Request or Cancel .	The query is canceled.
You start a query and close the browser.	The query is canceled.
You start a query and an system admin resets the IIS on the server.	The query is not canceled.
You start a query and leave your browser idle for longer than the session timeout specified in Relativity web.config, regardless of whether you clicked OK or Cancel on the message from the webpage pop-up.	The query is not canceled.
You start a query and click Cancel on the message from the webpage pop-up within the session timeout specified in Relativity web.config.	The query is canceled.
You start a query and click OK on the message from webpage pop-up within the session timeout specified in Relativity web.config.	The query continues to run. Relativity returns you back to the waiting screen (see the first row of this table).
You start a query and paste a different URL into your browser, or you refresh the page.	The query is canceled.

6.3 Frequently asked searching questions

This section includes frequently asked questions from Relativity users.

6.3.1 Multiple terms in dtSearch queries

Why would I receive an error message saying my query is too complex?

Relativity can't return precise results when a query includes search conditions that are too complex due to the use of multiple search terms. To avoid this error message, simplify your search criteria. For example, search for 10 instead of 50 terms in your query. You can save and tag your search results from each simpler query. With this approach you can search on the required terms, while providing better results.

6.3.2 Proximity searches in dtSearch queries

Why are terms in my proximity search highlighted even when they don't match my defined W/N criteria?

Relativity highlights terms that meet the requirements of your proximity search as well as the individual search terms. This behavior doesn't affect the results of your proximity search, which returns the appropriate documents. For example, the results of a proximity search for instances of law within three words of order (that is "law W/3 order") includes highlighted terms of "law" and "order" when they aren't within three words of each other. To accommodate this behavior, focus only on the results that match the proximity search criteria, and disregard the other highlighted terms.

6.3.3 Using reserved characters in dtSearch queries

Can I use a reserved character such as ampersand (&) in my dtSearch queries?

The following characters have special meaning in dtSearch, and must be treated with care when they appear in your queries: ?, *, %, ~, #, =, :, &, (). Use the dictionary to explore syntax options if necessary.

If you want to make any of these characters searchable, you must add them to the alphabet file and, in some cases, use Regular Expressions in your search. For more information, see the Searching Guide.

6.3.4 Multiple conditions

Why do my searches run slowly when I use multiple conditions?

When you add multiple search conditions to a query, Relativity searches on these conditions relative to each other, which slows down the return of your results. For example, you experience slow performance when running a query on all email messages received "after June 1" and "before June 30" of the same year. You can improve performance by using as few conditions as possible, such as excluding the condition "before June 30". Run the query with only the condition "after June 1", and then sort or filter your results to display messages received between the desired dates.

6.3.5 Nesting searches

Can I nest multiple searches in a saved search?

For performance reasons, we don't recommend nesting multiple searches in a saved search. You can select a search as a condition, but using multiple searches as conditions slows down the return of your results. See the Searching guide for more information.

7 Multiple object searching

This topic describes the logic behind multiple object searching.


When searching using *These Conditions* or *NOT These Conditions*, keep in mind the following:

- **These Conditions**—returns a document if any of the Relativity Dynamic Objects (RDO) that document is associated with matches the search criteria.
- **NOT These Conditions**—returns all documents that are not associated with at least one RDO that meets the criteria.

If a document is not associated with the object you're searching, it will not be returned in a search for *These Conditions*. For example, Document 3 below will never be returned in a search for *These Conditions* because it is not associated with any objects.

The image below shows documents associated with RDOs.


- Document 1 is associated with RDO1.
- Document 2 is associated with RDO2 and RDO3.
- Document 3 is not associated with any RDOs.
- Document 4 is associated with RDO4.



7.1 These Conditions

The following example returns Document 1 because:

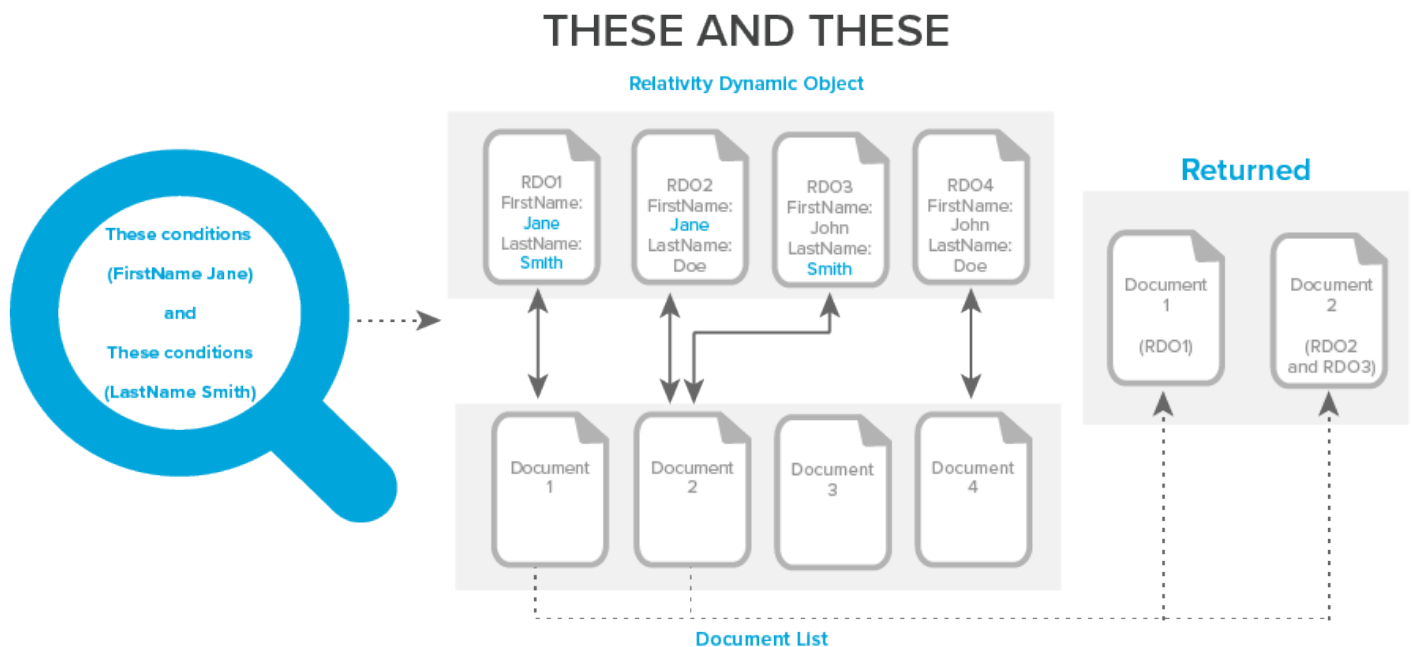
- Document 1 is associated with RDO1.
- RDO1 matches the search criteria (FirstName Jane, LastName Smith).



7.2 These Conditions and These Conditions

The following example returns Document 1 and Document 2 because:


- Document 1 is associated with RDO1.
- RDO1 matches the search criteria (FirstName Jane) and (LastName Smith).
- Document 2 is associated with RDO2 and RDO3.
- RDO2 matches the search criteria (FirstName Jane).
- RDO3 matches the search criteria (LastName Smith).



7.3 NOT These Conditions

The following example returns Document 3 and Document 4 because:

- Document 4 is associated with RDO4. RDO4 doesn't match the search criteria.
- Document 3 is not associated with any RDOs. Therefore, it does not match the search criteria.




7.4 These NOT

The following example returns Document 2 and Document 4 because:

- Document 2 is associated with RDO2 and RDO3.
- Although RDO2 contains (FirstName Jane), RDO3 does not (FirstName John). Remember, if one RDO does not meet the criteria, the document is returned.
- Document 3 is not returned because it isn't associated with any RDOs. Remember, you must associate a document with an object in order for it to return in a *These Conditions* search.

THESE NOT



7.5 None of these and Not all of these operators

The difference between *None of these* and *Not all of these* is subtle, but useful:

- **None of these**—will check the field and see if any of the choices on it match any of the conditions. If there are any choices checked on the field and one of them matches one of the conditions, the document will be rejected and not pulled by this search.
- **Not all of these**—will check the field and see if the choices on it match all the conditions. If yes, the document will be rejected. If only some of the choices match the conditions, the document will be pulled by this search.

Comparison of operators table

		Document 1 Coded Choice 1	Document 2 Coded Choice 1 and Choice 2	Document 3 No coding	Document 4 Coded Choice 2
Any of these	Choice 1	√	√		
	Choice 2		√		√
	Choice 1 and 2	√	√		√
None of these	Choice 1			√	√
	Choice 2	√		√	
	Choice 1 and 2			√	

		Document 1 Coded Choice 1	Document 2 Coded Choice 1 and Choice 2	Document 3 No coding	Document 4 Coded Choice 2
All of these	Choice 1	√	√		
	Choice 2		√		√
	Choice 1 and 2		√		
Not all of these	Choice 1			√	√
	Choice 2	√		√	
	Choice 1 and 2	√		√	√
	Is set	√	√		√
		√	√		√
		√	√		√
	Is not set			√	
				√	
				√	

8 Searching with the Entity object and name normalization

This topic describes the types of searches you can conduct using the Entity object and Name Normalization results.

8.1 Considerations

Note the following special considerations:

- The quality of your searches is affected by the quality of your name normalization results. Please use the [Name Normalization Quick Reference Guide](#) to verify your Name Normalization results.
- You can only create the following searches after running the name normalization operation.

8.2 Emails within a specific organization

To find emails between individuals within a specific organization, use the **Alias From** and **Alias Recipient** fields along with the associated **Alias From::Domain** and **Alias Recipient::Domain** fields. You must conduct this search with a double negative in order to find documents that only include your selected organization.

Notes:

- You can use variations of this search to identify specific cross organization communications. You can also use variations of this search to identify communications with outside counsel if you know the specific domain of outside counsel.
- There are scenarios where an alias doesn't include a domain (ex. Proper names). In these scenarios, the following search will not include documents where those aliases are communication. Reach out to Relativity Support for alternative workflows to circumvent this issue.

Create a search with the following conditions:

Note: Connect the search conditions as follows using Boolean operators: **1 AND 2 AND 3**

Condition 1

- **Field:** Alias From
- **Operator:** is set

Condition 2

- **Field:** Alias From
- **Operator:**not these conditions
 - **Field:** Alias From::Domain
 - **Operator:** is not
 - **Text:** [Domain 1]

Condition 3

- **Field:** Alias Recipient
- **Operator:** not these conditions
 - **Field:** Alias Recipient::Domain
 - **Operator:** is not
 - **Text:** [Domain 1]

1 AND 2 AND 3

1. Alias From
these conditions

Alias From
is set

AND ▼

2. Alias From
not these conditions

Alias From::Domain
is not enron.com

AND ▼

3. Alias Recipient
not these conditions

Alias Recipient::Domain
is not enron.com

Breaking it down:

- **Alias From::Domain==not this domain** means someone without this domain is the sender
- **Not these conditions [LINE ABOVE]** means no domain other than this domain is the sender

8.3 Communications between two specific individuals and no one else

To find emails between two individuals and no one else, use the **Entity From** and **Entity Recipient** fields. You must conduct this search with a double negative in order to find documents where only your selected entities appear.

Notes:

- You can use variations of this search to identify communications where a specific entity is the only recipient, but the senders vary.
 - You can also use variations of this search to exclude CC or BCC values.
 - You can use the Entity Participant field to ensure no one else existed on the email thread.
-

Create a search with the following conditions:

Note: Connect the search conditions as follows using Boolean operators and Logic Groups: **1 AND ((2 AND 3) OR (4 AND 5))**

Condition 1

- **Field:** Entity From
- **Operator:** is set

Condition 2

- **Field:** Entity From
- **Operator:**not these conditions
 - **Field:** Entity From
 - **Operator:** none of these
 - **Text:** [Entity A]

Note: This is a double-negative condition that will **exclude** everyone but Entity A. In other words, this condition will only allow emails **sent by Entity A**. When conditions 2 and 3 are combined, we have a holistic search condition for emails sent by Entity A and received by Entity B.

Condition 3

- **Field:** Entity Recipient
- **Operator:**not these conditions
 - **Field:** Entity Recipient
 - **Operator:** none of these
 - **Text:** [Entity B]

Note: This is a double-negative condition that will **exclude** everyone but Entity B. In other words, this condition will only allow emails **received by Entity B**. When conditions 2 and 3 are combined, we have a holistic search condition for emails sent by Entity A and received by Entity B.

Condition 4

- **Field:** Entity From
- **Operator:**not these conditions
 - **Field:** Entity From
 - **Operator:** none of these
 - **Text:** [Entity B]

Note: This is a double-negative condition that will **exclude** everyone but Entity B. In other words, this condition will only allow emails **sent by Entity B**. When conditions 4 and 5 are combined, we have a holistic search condition for emails sent by Entity B and received by Entity A.

Condition 5

- **Field:** Entity Recipient
- **Operator:** not these conditions
 - **Field:** Entity Recipient
 - **Operator:** none of these
 - **Text:** [Entity A]

Note: This is a double-negative condition that will **exclude** everyone but Entity A. In other words, this condition will only allow emails **received by Entity A**. When conditions 4 and 5 are combined, we have a holistic search condition for emails sent by Entity B and received by Entity A.

1 AND ((2 AND 3) OR (4 AND 5))

1. Entity From

these conditions

Entity From

is set

AND ▼

2. Entity From

not these conditions

Entity From

none of these: [Wayne](#), [Bruce](#)

AND ▼

3. Entity Recipient

not these conditions

Entity Recipient

none of these: [Parker](#), [Peter](#)

OR ▼

4. Entity From

not these conditions

Entity From

none of these: [Parker](#), [Peter](#)

AND ▼

5. Entity Recipient

not these conditions

Entity Recipient

none of these: [Wayne](#), [Bruce](#)

Note: By combining the two logic groups above, we create a holistic set of search conditions that will return all emails **sent from entity A to entity B** as well as all emails **sent from entity B to entity A**.

Breaking it down:

- **Entity From==not this person** means someone other than this person is the sender
- **Not these conditions [LINE ABOVE]** means nobody other than this person is the sender


8.4 Communications between any two individuals and no one else

To find emails between any two individuals and no one else, use the **Email Recipient Count** field created by Relativity Processing or the **Delimiter Count by Saved Search** script to count the number of recipients for each document.

Note: If your workspace does not have the **Delimiter Count by Saved Search** script, download and install it. See [NSerio / delimiter-count-by-saved-search](#) on GitHub for a link to the script. For more information on adding a script from the library, see Scripts documentation.

Create the following search:

- **Field:** Email Recipient Count
- **Operator:** is
- **Text:** 1



Because an email always has a single sender, this search focuses on finding emails that have just one recipient.

Note: You can also use this search can to filter out email blasts by looking for documents where **Email Recipient Count** is greater than 20.

8.5 Emails across entity metadata

You can use any entity metadata field for advanced searching on documents. As an example, to find emails sent from the Finance department to the Human Resource department, you can use the **Entity From::Department** and **Entity Recipient::Department** fields.

Notes:

- You can use variations of this metadata search to identify conversations with attorneys, job title, between your Chicago and Hong Kong offices, location, or between males and females, gender, and more.
- You may need to edit the Entity field to set **Open to Associations** to **Yes** in order to search across it on the document object.
- Importing address book, Active Directory, or CRM information for each entity can give you more metadata to leverage when searching across documents.

Create a search with the following conditions:

Note: Connect the search conditions as follows using a Boolean operator: **1 AND 2**

Condition 1

This search indicates that a specific entity value is the sender of the document.

- **Field:** Entity From
- **Operator:** these conditions
 - **Field:** Entity From::Department
 - **Operator:** is
 - **Text:** [Department 1]

Condition 2

- **Field:** Entity Recipient
- **Operator:** these conditions
 - **Field:** Entity Recipient::Department
 - **Operator:** is
 - **Text:** [Department 2]

The screenshot shows a search builder interface with two conditions connected by an AND operator. The first condition is '1. Entity From' with the text 'these conditions' and 'Entity From::Department is Finance'. The second condition is '2. Entity Recipient' with the text 'these conditions' and 'Entity Recipient::Department is Human Resources'. The conditions are connected by an 'AND' operator with a dropdown arrow.

1 AND 2

1. Entity From
these conditions
Entity From::Department
is Finance

AND ▼

2. Entity Recipient
these conditions
Entity Recipient::Department
is Human Resources

8.6 Emails sent from two different internal domains

To find emails sent between two different internal domains, use the **Alias From** and **Alias Recipient** fields. Use the **Text** field to define each domain. The logic below pulls emails sent from Domain 1 to Domains 1 or 2 and emails sent from Domain 2 to Domains 1 or 2.

Logic Group 1

- **Condition 1**
 - **Field:** Alias From
 - **Operator:** these conditions
 - **Operator:** is
 - **Text:** [Domain 1]
- **Sub-logic Group 1B**
 - **Condition 2**
 - **Field:** Alias Recipient
 - **Operator:** not these conditions
 - **Field:** Alias From::Domain
 - **Operator:** is not
 - **Text:** [Domain 1], [Domain 2]
 - **Condition 3**
 - **Field:** Alias Recipient
 - **Operator:** these conditions
 - **Operator:** is set

Logic Group 2

- **Condition 4**
 - **Field:** Alias From
 - **Operator:** these conditions
 - **Operator:** is
 - **Text:** [Domain 2]
- **Sub-logic Group 2B**
 - **Condition 5**
 - **Field:** Alias Recipient
 - **Operator:** not these conditions
 - **Field:** Alias From::Domain
 - **Operator:** is not
 - **Text:** [Domain 1], [Domain 2]

- **Condition 6**
 - **Field:** Alias Recipient
 - **Operator:** these conditions
 - **Operator:** is set

8.7 Emails someone sent to themselves

To find emails someone sent to themselves and nobody else, use the “Entity From” and “Entity Recipient” fields. You must conduct this search with a double negative in order to find documents that only include our selected entity.

Note: This search may be impacted by whether you decided to group professional and personal aliases under a single entity or separate entities.

Create a search with the following conditions:

Note: Connect the search conditions as follows using a Boolean operator: **1 AND 2 AND 3**

Condition 1

- **Field:** Entity From
- **Operator:** these conditions
 - **Field:** Entity From
 - **Operator:** is set

Condition 2

This search indicates that someone other than this person is the sender.

- **Field:** Entity From
- **Operator:** not these conditions
 - **Field:** Entity From
 - **Operator:** none of these
 - **Text:** [Entity 1]

Search 3

- **Field:** Entity Recipient
- **Operator:** not these conditions
 - **Field:** Entity Recipient
 - **Operator:** none of these
 - **Text:** [Entity 1]

1 AND 2 AND 3

1. Entity From
these conditions

Entity From
is set

AND ▼

2. Entity From
not these conditions

Entity From
none of these: [Wayne, Bruce](#)

AND ▼

3. Entity Recipient
not these conditions

Entity Recipient
none of these: [Wayne, Bruce](#)

Breaking it down:

- **Entity From==not this person** means someone other than this person is the sender
- **Not these conditions [LINE ABOVE]** means nobody other than this person is the sender

8.8 Entities that communicated on privileged documents

To find a list of entities that communicated on privileged documents, use the **Entity Participant::Privilege Designation** field on the Entity object (Entities tab).

Notes:

- You must edit the Privilege Designation field to set **Open to Associations** to **Yes** in order to search across it on the document object.
- You can use a variation of this search using the Participant field to return the entities that only sent or received privileged documents rather than participated in any segment of the document.
- You can use variations of this search to return the entities that communicated on a document that have any specific attribute. For example, entities that communicated on Hot documents.

Create the following search:

- **Field:** Entity Participant
- **Operator:** these conditions
 - **Field:** Entity Participant::Privilege Designation
 - **Operator:** any of these
 - **Text:** [Designation Type(s)]

This search locates Entities that exist in the **Entity Participant** field of documents that have this specific value in the **Privilege Designation** field.



8.9 Emails where a specific person drops off of a communication

To find emails where a specific person dropped off an email chain, use the **Entity From**, **Entity Recipient**, and **Entity Participant** fields. This search finds all documents where the person appears in a lower email segment within the document (Entity Participant), but not the top email segment (Entity From and Entity Recipient).

Note: Adding multiple people to the search can cause documents to be missed. If you want to search for multiple people, create separate search cards for each person.

Create a search with the following conditions:

Note: Connect the search conditions as follows using a Boolean operator: **1 AND 2 AND 3**

Condition 1

Entity Participant (these conditions), Entity Participant (any of these [Entity 1])

- **Field:** Entity Participant
- **Operator:** these conditions
 - **Field:** Entity Participant
 - **Operator:** any of these
 - **Text:** [Entity]

Condition 2

Entity From (not these conditions), Entity From (any of these [Entity 1])

Condition 3

Entity Recipient (not these conditions), Entity Recipient (any of these [Entity 1])

1 AND 2 AND 3

1. Entity Participant

these conditions

Entity Participant

any of these: [Parker](#), [Peter](#)

AND ▼

2. Entity From

not these conditions

Entity From

any of these: [Parker](#), [Peter](#)

AND ▼

3. Entity Recipient

not these conditions

Entity Recipient

any of these: [Parker](#), [Peter](#)

9 dtSearch

Relativity's dtSearch engine provides advanced search functionality such as proximity, stemming, and fuzzy searches across any field type. It also supports the use of Boolean operators and custom noise word lists and the basic searching features available in keyword searches. After building your dtSearch index, the Dictionary search option becomes available.

Note: Relativity partitions a single index into smaller indexes, called sub-indexes, which multiple workers build simultaneously. This increases performance by spreading out the work over a configurable number of agents. When you perform a search, Relativity runs your query on the smaller indexes in parallel. The application then federates and returns your results. For more details, see the Ask The Expert Training content: [Searching: Best Practices for dtSearch Builds](#).

9.1 dtSearch index

You can build custom dtSearch indexes for a subset of documents or for certain document fields in a workspace. You must have the appropriate permissions to complete this task. See [Workspace security](#) in the Admin Guide.

Before you begin, you need to create a saved search that includes the fields that you want to include in the index. You can then name the index based on the document search set used to create it.

Note: Within a field, dtSearch truncates any string longer than 32 characters that does not contain a space character. It indexes only the first 32 characters of the string. For more information, see [Searching for words longer than 32 characters](#).

9.1.1 Accent-insensitive indexes

By default, Relativity builds an accent-insensitive index. In an accent-insensitive index, some characters translate to the base character, which causes those characters and any terms containing those characters to be treated the same in a search terms report.

Note: dtSearch uses .ABC files, but only for characters in the range from 33-127. Relativity handles all other characters according to the definitions in the Unicode character tables.

When searching for characters in a character-insensitive index, accented characters are typically treated as their base, unaccented form. This simplifies searching because it does not matter if a character has an accent or not.

What to expect:

- Accented characters will show up as their base character. For example, “ṽ” will be shown as “v”.
- The search will not show you the accented character, only the base character.
- You will not see the breakdown, decomposition, of how a character is split into simpler parts. You can still find the right base character.
- **Example:** Relativity converts accented characters like á or ñ to the unaccented versions, a or n.
- **Example:** If you search for the term fröhlich, searching that term as fröhlich or frohlich would both return the hit. However, highlighting in the Viewer may not display both variations.

If you want more detailed information about how characters are broken down or mapped, you can use tools like the [Compart Unicode tool](#) for extra help. The way that Compart shows that a character is a base character is by omitting the decomposition row entirely. For example, the letter [v on the Compart website](#).

9.1.2 Creating a dtSearch index

To create a new dtSearch index:

1. Use the search feature to navigate to the **Search Indexes** tab.
2. Click **New Search Index**.
The dtSearch Index Information form appears with required fields having an orange asterisks.
3. Complete the fields on the dtSearch index form. See more details, see [Fields](#).
4. Click **Save** to display the index details page. The index details page now displays three additional read-only fields and the dtSearch index console. See [Fields](#) and [dtSearch console](#).
5. Click **Build Index: Full**. A dialog box asks you to verify that you want to run a full build. You can also select **Activate this index upon completion**. Indexes must be active to search them.
6. Click OK to build your index.

Note: Network problems can slow down your dtSearch builds. If a dtSearch manager or worker agent encounters a network-related error during the build process, it executes up to three retry attempts at 30-second intervals.

7. If you did not select **Activate this index upon completion** in the dialog box, click **Activate Index** on the console. The index will not activate if there are errors. Activating an index makes it available in the Search menu.
8. (Optional) Click **Refresh Page** at any point in the build to see the index's current build status. If errors occur during the build, the **Retry Errors** button enables on the console under the Errors and Status heading. Click this button to try to resolve any errors.

Once the index builds, the console enables more options. See [dtSearch console](#).

9.1.2.1 Fields

The dtSearch index page includes the following fields:

dtSearch Index Information

Add information to the following dtSearch index fields:

- **Name**—the dtSearch index name. This name appears within the **search with** menu in the Documents tab.
- **Order**—the integer value, positive or negative, representing the position of the index in the search indexes list. Indexes sort from lowest, top, to highest, bottom, order number. Those with the same order number sort alpha-numerically.
- **Searchable set**—the saved search of documents for indexing. Relativity indexes the documents returned by the search as well as the returned documents' fields. It may use a dtSearch or an Analytics index. Make sure the index is active.

Note: When creating a saved search for a dtSearch index, the best practice is to use long text fields. You will see a warning message, when creating a dtSearch index, if you select a saved search that does not contain at least one long text field.

- **Email notification recipients**—specifies recipients to send an email notification to when your dtSearch index finishes running. Enter the email addresses of the recipients. Separate entries with a semicolon.
- **Notes**—enter any considerations, or thoughts, about the search index.

Searchable Set - Current Indexable Fields

This section is a list of the fields and field types in the saved search used for indexing.

Note: The fields listed may not match those indexed if someone made updates after indexing the searchable set.

If a long-text field not included in the searchable set, it's possible to not have search any search results returned. For better search results, check the saved search to ensure a field with the field type as Long Text. For example, extracted text is included.

9.1.2.2 Advanced Settings

Add information to the following advanced setting fields:

- **Auto recognize date, email, and credit card numbers**—a yes/no field. See [Auto-recognition](#) for details.
- **Create accent sensitive**—a yes/no field. When set to Yes, dtSearch indexes are sensitive to accents and other language-specific characters.
- **Skip malicious files**—a yes/no field. Select **Yes** to skip potentially malicious quarantined files and continue processing the index. The final index breakdown includes the number of skipped documents.
- **Use Default Sub-index Size**—select **Yes** to use the default sub-index size of 16 GB. Sub-index size is automatically optimized for system performance, unless you select **No**. Then you must enter sub-index size.
- **Custom Sub-Index Size (GB)**—determines the size of each sub-index created when you generate a dtSearch index. The field must contain a numerical value between 1-16 GBs.
- **Index share**—the location of the file share for storing the search index. Select the bold (default) location. You will see a warning about system performance if you select a location other than the default.

9.1.2.3 Noise Words

Edit the list of words that Relativity ignores during indexing.

9.1.2.4 Alphabet

Edit the index's alphabet file. See [Making a character searchable](#).

Note: If you search for long, uninterrupted strings that have no spaces or word breaks, such as when you have made a character searchable, dtSearch truncates the string after 32 characters and inserts a wildcard. For more information, see [Searching for words longer than 32 characters](#).

9.1.3 dtSearch index page

After you create and build a dtSearch index, the dtSearch page has several sections where you can view details about your index.

9.1.3.1 Index Status

You can view the state of your dtSearch index from the Index Status section. The name of the Index Status section populates with the name of your dtSearch index. When you are building an index, this section changes to a progress bar where you can track your index's progress in real-time. When the index is no longer in progress, this section changes to a static field that displays the below fields.

- **Status**—the status of the index. For example, Active - Indexed or Inactive - Indexed.
- **Document Breakdown** —the number of indexed documents.

9.1.3.2 dtSearch Index Information

The dtSearch Index Information section provides general details about the settings applied to your dtSearch index. This section has the following information:

- **Name**—the name of your index.
- **Order**—the integer value, positive or negative, representing the position of the index in the search indexes list. Indexes sort from lowest, top, to highest, bottom, order number. Those with the same order number sort alpha-numerically.
- **Searchable set**—the set of documents to be indexed. You can choose from any saved search in the workspace.
- **Email notification recipients**—the emails that receive an email notification when your index population fails or completes.
- **Notes**—the notes added by the user that created the index.

9.1.3.3 Searchable Set - Current Indexable Fields

The fields listed may not match those indexed if someone made updates after indexing the searchable set.

9.1.3.4 Advanced Settings

The Advanced Settings section provides sub-index details about your dtSearch index. This section has the following information:

- **Auto-recognize date, email, and credit card numbers**—a yes/no field.
- **Create accent sensitive**—a yes/no field. When set to Yes, dtSearch indexes are sensitive to accents and other language-specific characters.
- **Skip malicious files**—a yes/no field. Select Yes to skip potentially malicious quarantined files and continue processing the index. The final index breakdown includes the number of skipped documents.
- **Use Default Sub-Index Size**—a yes/no field that determines whether to use the default 16 GB size or a custom size. *Custom Sub-Index Size* field.
- **Index share**—the location of the file share for storing the search index. Select the bold (default) location. You will see a warning about system performance if you select a location other than the default.

9.1.3.5 Noise Words

The list of words that Relativity ignored during indexing.

9.1.3.6 Alphabet

The index's alphabet file. See [Making a character searchable](#).

Note: If you search for long, uninterrupted strings that have no spaces or word breaks, such as when you have made a character searchable, dtSearch truncates the string after 32 characters and inserts a wildcard. For more information, see [Searching for words longer than 32 characters](#).

9.1.3.7 Temporary Index Details


The Temporary Index Details section only appears during an incremental build. This table displays sub-indexes that copy from your original index and are in the process of modification during the incremental build. Once the sub-indexes in this table update, they replace the original sub-indexes from which they were copied. This section has the following information:

- **Population Table**—the name of the table that a sub-index is populating.
- **Build Status**—the state that the sub-index is currently in.
- **Worker Agent**—the name of the agent that is handling the sub-index.
- **Worker Agent Status**—the current state of the worker agent.
- **Index File Share**—the location where you store your sub-index.
- **Document count**—the number of documents assigned to the sub-index.
- **Error(s)**—any errors encountered by the sub-index.
- **Fragmentation Level**—the fragmentation level of the sub-index. Any index equal to or greater than the sub-index fragmentation threshold appears in red.

9.1.3.8 Current Index Details

The Current Index Details section displays the sub-indexes that make up your dtSearch index. This section has the following information:

- **Population Table**—the name of the table that a sub-index is populating.
- **Build Status**—the current state of the sub-index.
- **Worker Agent**—the name of the agent that is handling the sub-index.
- **Worker Agent Status**—the current state of the worker agent.
- **Index File Share**—the location where you store your sub-index.
- **Document count**—the number of documents assigned to the sub-index.
- **Error(s)**—any errors encountered by the sub-index.
- **Fragmentation Level**—the fragmentation level of the sub-index. Any index equal to or greater than the sub-index fragmentation threshold appears in red.
- **Rebuild Selected Sub-Indexes**—manually rebuilds selected sub-indexes. Do not use this option unless directed by the [Support](#) team.



9.1.3.9 View Audit

Using the View Audit button, you can see when a user modified dtSearch index settings. The View Audit layout has the following fields:

- Action
- Field
- Old Value
- New Value
- User Name

- Timestamp
- Details

9.1.3.10 Temporary storage

If you specify a temporary storage location, dtSearch builds the index in this directory and then copies the index over to the final index share when the build completes. Using a temporary storage location could speed up the build time and reduce network contention

9.1.4 dtSearch console

The dtSearch index console includes the following options:

- **Build Index: Full**—creates a full build of the index. During the build, the button toggles to **Cancel Build**.

Note: You must perform a full build when you add an additional field to the index, change any index settings, change fields of the searchable set, or overlay text on existing fields.

- **Cancel build**—canceling the build stops the indexing thread, leaving the index in an unstable state.
 - Relativity deletes these indexes from the population table and gives them an inactive status.
 - You cannot search on an index with an inactive status until you run a full build.
 - Canceling also deletes the index files in the index share.
- **Build Index: Incremental**—updates an index after adding or removing documents.
 - Incremental builds do not modify the text of any documents included in previous builds, even if the original fields are no longer returned in the source search.

Note: If you are modifying the text of already indexed documents, you must either run a full build or re-index documents with an incremental build. If you do not want to perform a full build, you can follow the [Re-index documents with incremental builds on the next page](#) workflow to remove the impacted documents from the index. Then add them back with the new text.

 - During an incremental build the existing index remains available for searching, but changes to the index are not reflected in search results until the incremental build completes.
 - Canceling an incremental build returns the index to its previous state.
 - The incremental build process copies each sub-index that requires modification, updates the copy, then replaces existing sub-indexes with the updated copies.
 - When run, the Case manager agent removes duplicate sub-indexes.
 - The system automatically compresses a sub-index during an incremental build only if the sub-index fragmentation level is equal to or greater than the sub-index fragmentation threshold value.
- **Compress Index**—compresses the dtSearch index returning all sub-indexes with a fragmentation level greater than zero to a fragmentation level of zero.
 - The Compress Index button only runs compression on sub-indexes that have a fragmentation level greater than zero.
 - You can search on the original, uncompressed, dtSearch index while compression is in progress.
 - Once compression completes, the system automatically replaces the old sub-indexes with the defragmented sub-indexes.

- When run, the Case manager agent removes duplicate sub-indexes.
- Canceling compression returns the index to its original fragmented state before compression began.
- **Deactivate Index**—deactivates the index and removes it from the **search with** menu in the Documents tab, but not from the database.
- **Swap Index**—swaps your index with a replacement index to use its resources while your index builds or is inactive or disabled for any reason.
 - This enables you to keep searching while your primary index experiences downtime.
 - The Swap Index function updates anything in the Views table, which affects batches, saved searches, nested searches, and more.
 - You can only select indexes in the Replacement Index with an Active status.
 - This index you swap to does not automatically run an incremental update.
 - Selecting the index from the drop-down list and clicking **OK** completes the index swap.
 - You cannot reverse the swap results in the current dialog box.
 - You must close this swap and run it again to swap back or swap another time.
 - This functionality is useful in limited cases. For example, if you are performing a full rebuild on a very large index.
 - Since dtSearch incremental builds are online, you can search documents once indexed.
- **Retry Errors**—enables only if errors occur, you can use this button to resolve errors.
- **Show Document Errors**—enables only if document errors occur. This button creates an exportable list of document-level errors.
- **Show Detailed Status**—shows you statistical data for the index, including:
 - **Doc Count**—the total number of documents in the index.
 - **Index Size**—the size of the index in bytes.
 - **Created Date**—the date you created the index.
 - **Updated Date**—the date you updated the index.
 - **Last Build Duration**—how long the last build took to complete in hours, minutes, and seconds.
- **Refresh Page**—shows the index's current build status.

9.1.5 Re-index documents with incremental builds

In large workspaces, dtSearch indexes can take a long time to build. This results in users being unable to run dtSearch search terms. In these situations, incremental builds are favored over full builds because users can run searches while building. However, incremental builds only add or remove documents. Incremental builds do not update the text of existing documents. The following workflow can be used as a workaround.

Due to Relativity's structure of sub-indexes, the following workflow will re-index existing documents to bring in updated extracted text or other modified fields for already indexed documents:

1. Change the dtSearch index's searchable set to a search with the documents you want to keep. Such as removing the ones you want to re-index.
2. Perform an incremental build. Incremental builds remove documents no longer in the searchable set as of Search Grid's implementation.

3. Perform an incremental build and those records will now be indexed again.
4. Perform an incremental build to index the new documents that were added back to the search.

Note: The dtSearch index will be active and searchable during the incremental builds. This means search terms will not contain hits for the documents that are being re-indexed until step 4 is completed.

9.2 Making the dtSearch noise word list searchable

Relativity ignores words that do not act as meaningful criteria when you create dtSearch and keyword queries. Ignored words are known as noise or noise words. Search indexes automatically include the default list of noise words. However, you can edit this list in the dtSearch list to suit your needs. This article includes an overview of noise words and steps to create custom lists.

9.2.1 Default noise word list

Relativity references the default list of noise words each time you create a new index. System admins cannot edit noise words in keyword searches. The default noise word list consists of punctuation marks, single letters and numbers, and the following words:

Begins with	Noise words
A	a, about, after, all, also, an, and, another, any, are, as, at
B	be, because, been, before, being, between, both, but, by
C	came, can, come, could
D	did, do
E	each, even
F	for, from, further, furthermore
G	get, got
H	had, has, have, he, her, here, hi, him, himself, his, how, however
I	i, if, in, indeed, into, is, it, its
J	just
L	like
M	made, many, me, might, more, moreover, most, much, must, my
N	never, not, now
O	of, on, only, or, other, our, out, over
S	said, same, see, she, should, since, some, still, such
T	take, than, that, the, their, them, then, there, therefore, these, they, this, those, through, thus, to, too
U	under, up
V	very
W	was, way, we, well, were, what, when, where, which, while, who, will, with, would
Y	you, your

Note: Relativity ignores noise words. However, Relativity does not ignore their position in the search phrase set. So, if you execute the query apple w/6 pear, the search returns the phrase apple tree is far from the pear even though it contains the noise words is, from and the.

9.2.2 dtSearches and noise words

The default list of noise words is the same in a dtSearch as in a keyword search. The primary difference is that you can customize the dtSearch index list. For example, if the word never is important to your litigation, remove it from the noise words list, so that your search results always return that word.

To create a custom noise word list, perform the following:

1. Create a new dtSearch index, and then name it *dtSearch - updated noise words*.
2. Select your extracted text search for the Searchable set.
3. Delete the word *never* from the Noise Words list.
4. Save the list, and then perform a full build on your new index.


9.2.3 Noise words in languages other than English

You can set up noise words to search documents in other languages. If the workspace primarily contains of documents in a different language, see for an overview of suggested noise words for use in nineteen additional languages.

9.3 Alphabet list

Some of the characters in the alphabet file are not printable. Screenshots were used instead of the actual text. You cannot copy or paste the **Spaces** or **Ignore** characters since they are not printable. Instead, use the *dtSearchDefaultAlphabetFile* instance setting in the Instance Setting system guide to update the dtSearch default alphabet file.

Note: Each sequence must start with a leading, or empty, space. Not having the leading space may produce errors.



9.3.1 dtSearch Alphabet File

The following is the default dtSearch Alphabet file you'll find in Relativity. It contains letters, numbers, hyphens, spaces, ignore, and end (CJK ranges) sections.

[Click here to view the default dtSearch Alphabet file.](#)

[Letters] // Original letter, lower case, upper case, unaccented

```
0000  
1111  
2222  
3333  
4444  
5555  
6666
```

7777
8888
9999
AaAA
BbBB
CcCC
DdDD
EeEE
FfFF
GgGG
HhHH
IiII
JjJJ
KkKK
LILL
MmMM
NnNN
OoOO
PpPP
QqQQ
RrRR
SsSS
TtTT
UuUU
VvVV
WwWW
XxXX
YyYY
ZzZZ

-- --
aaAa
bbBb
ccCc
ddDd
eeEe
ffFf
ggGg
hhHh
iili
jjJj
kkKk
lIlI
mmMm
nnNn
ooOo
ppPp
qqQq
rrRr
ssSs
ttTt
uuUu
vvVv
wwWw
xxXx
yyYy
zzZz

[Hyphens]

-

[Spaces]
\\09\\0a\\0c\\0d !@\"#\$%&'()*+,-./:;<=>?[\\5c]^`{|}~

[Ignore]
00000000\\080000000000000000000000000000%

[End]
CJKRanges = 0e00-0e4e 3040-30ff 4e00-9fff

Note: Previous guidance hid some characters that should not have been hidden. Those non-printable characters are critical to index function and should never be removed. The "\\08" text represents the "backspace" text, and should also never be removed or split up (that is, trying to index the backslash).

Note: This section only accepts ASCII characters (code points between 33 and 127) as input. We currently do not support ignoring extended Unicode values.

9.3.2 Alphabet file validation

When you save a dtSearch index, Relativity runs a validation check on the alphabet list. You will see a warning message if Relativity detects invalid spacing or syntax. You cannot save the index if there are errors with the alphabet list.

The validation check includes:

- Header sections
 - Header section appears first in Alphabet
 - Exact header section without any added whitespace
 - Required newline before section
- Letters
 - Exact title, allowing any whitespace and comments preceding double slash //
 - Each letter on own line with preceding space
 - Each letter variant separate by single space
 - Allow any extra whitespace after letter
- Hyphens, Spaces, and Ignore
 - Exact title, allowing any whitespace
 - Single line of characters with preceding space
 - Optional newlines before next section
- Footer sections
 - Exact title
 - Skip validating any text following title
- General
 - Purple, Pink, Red, Green sections are each optional and can be in any order

9.3.3 Alphabet file sections

The following descriptions are for characters in the ASCII 33-127 range.

9.3.3.1 Letters

dtSearch defines letters as characters to index. This includes all alphabetical characters (a-z and A-Z) and all digits (0-9).

Note: dtSearch is case insensitive. You cannot make dtSearch case-sensitive in Relativity by modifying the Letters section of the Alphabet file.

9.3.3.2 Hyphens

dtSearch defines hyphens as characters that receive special processing in dtSearch. By default, dtSearch only classifies the - character as a hyphen.

[Hyphens]

-

Searching for hyphens

Hyphens are separated into their own category because dtSearch has special processing for them. Hyphens in this section include the normal hyphen (-) the em dash (—) and the en dash (–). Relativity does not distinguish these as different, regardless of what goes into the hyphen section. By default, dtSearch will index the hyphen as a space.

In the following example, we have three documents:

Doc 1 = "I sent it by first-class mail"

Doc 2 = "I sent it by first class mail"

Doc 3 = "I sent it by firstclass mail"

A search for *first class* on an dtSearch index with default noise words and Alphabet file, will return documents 1 and 2.

If we remove the Hyphen from the hyphen section, then a search for *first class* will return documents 1 and 2.

To force dtSearch to recognize the hyphen characters in dtSearch, add it to the alphabet section of the Alphabet File. To do this, see [Searching for symbols and emojis on page 79](#), and replace all % signs with the hyphen (-) character.

Note: You will need to remove the hyphen from the Hyphen section, and it does not appear in the spaces section. You will not need to use a regular expression to search for the hyphen.

After running a full build, dtSearch will recognize the hyphen as a searchable character. This will return cases where all three commonly used hyphens are used: the hyphen, em dash and the en dash. There is no way to distinguish between the hyphen, em dash, and the en dash.

After making this change, a search for *first class* will only return Doc 2. A search for *first-class* will only return Doc 1, and a search for *firstclass* will only return Doc 3.

Here is the dtSearch Hyphen page: <https://support.dtsearch.com/dts0154.htm>

9.3.3.3 Spaces

dtSearch defines a space character as a character that causes a word break. These characters are not indexed and are not searchable. By default, dtSearch treats the following characters as spaces:

[Spaces]

\09\0a\0c\0d !@\"#\$%&'()*+,-./:;<=>?[\5c]^`{|}~

Values listed as \## are Unicode code points. Their definitions are:

- \09—horizontal tab
- \0a—line feed

- \0c—form feed
- \0d—carriage return
- \5c—backslash (\)

For more information, log into the Relativity Community to review the [dtSearch Unicode values for Special Characters](#) article.

Searching for an asterisk

You can do so by using a Regular Expressions (RegEx) workflow to get around the fact that * is a reserved dtSearch operator. Wrapping up (or you can think of this as insulating) the asterisk * in a RegEx allows * to be treated as the literal character *, and not the wildcard operator *.

You will need to do the following:

- Make the asterisk * a searchable character in your dtSearch index by adjust the alphabet file. Building a second dtSearch index is recommended, as opposed to making * searchable in the main dtsearch index.
 - Refer to the recipe “Regular expression searching – symbols“, which contains detailed instructions on how to make characters searchable . In the recipe, substitute the % sign for the asterisk *.
- Use Regular Expressions (RegEx) to bring back the proper results.

Now, the recipe directs you to use the RegEx metacharacter \W to search for the % sign. Instead, you will just need to escape the * asterisk with a backslash \. Once you’ve made * searchable, the exact RegEx search terms you will need to plug into your dtSearch box to search for **REHAB is:

```
"##\*\*rehab"
```

"##" signals to dtSearch that what’s contained in between needs to be treated as a RegEx. The backslash \ escapes the * character so that it is treated as * and not a RegEx metacharacter.

In this example below, I’m searching for **REHAB and have made * a searchable character in my index. Note first the test data. The goal is to bring back only those records on lines 1 and 2.

When searching for **REHAB using the RegEx "##**rehab", I get the correct results:

Rehab by itself will not work even though you made * a searchable character because without the RegEx, * is still treated as the wildcard operator. Surrounding "*Rehab" in quotes also will not work.

Search for parentheses

Parenthesis are used for grouping search request terms in Boolean expressions. Quoting them makes them part of the search term instead. However it is possible to make the parentheses searchable by taking the following steps:

1. Add the parentheses to your alphabet file, using the steps in documentation for adding symbols. For more information, see [Searching for symbols and emojis on page 79](#)
2. Run a full build on your dtSearch index.
3. Add quotes around your dtSearch search term when searching. Quotes tell dtSearch we want the term in the quotes treated as characters, not syntax.

For example, if you are searching for the term (apple) with parenthesis in the term, then you would just add quotes around the term like so: "(apple)".

Permanently adding (and) to the alphabet file is not recommended because of the role of parentheses in search requests as logical operators, which will be confusing for users. Worse, parentheses in document text will cause search terms to be effectively lost. For example, consider this excerpt:

The defendant (John Smith) ...

If parentheses are indexed, then neither John nor Smith will be indexed, because the parentheses become part of the word. The search terms in the index would be "(John" and "Smith)". An index created this way is effectively useless for anything except specifically searching for parenthesis.

Because of this, we recommend that you only make parentheses searchable in a separate dtSearch index for use with the searches which contain the parentheses.

9.3.3.4 Ignore

dtSearch defines an ignored character as a character that is not indexed and does not create a word break when processing text. These characters are not searchable. By default, dtSearch ignores the following characters:

[Ignore]

```
00000000\080000000000000000000000%
```

- \08—backspace character
- %—percent sign

Values listed as ### are Unicode code points. Currently, the only code point that is ignored by default is the backspace character (\08).

9.3.3.5 End

dtSearch has defined ranges for CJK characters and these will make each Thai, Chinese, and Japanese character a separate word.

[End]

```
CJKRanges = 0e00-0e4e 3040-30ff 4e00-9fff
```

For more information, see the Searching guide.

9.3.3.6 Non-ASCII characters

Non-ASCII characters have a Unicode value greater than 0x7F. Many characters that are not ASCII are searchable by default. For those which are not, for example € and £, you can index them by adding their hexadecimal code to the AdditionalLetters section of the alphabet file. For more information, see [Searching for symbols on the next page](#).

9.3.3.7 Restricted characters

Some characters cannot be queried with standard syntax because of a limitation in dtSearch or because of how Relativity uses the dtSearch API. The following characters require special treatment in your query:

" () * ? % ~ # & =

For searching with parenthesis, see [Search for parentheses on the previous page](#)

You can use a regular expression to search for these characters. For example, [Searching for an asterisk on the previous page](#).

9.3.3.8 Searching for a symbol or character

To search for a symbol or character in Relativity, see [Searching for symbols on the next page](#).

9.3.3.9 Searching for emojis

To search for emojis in Relativity, see [Searching for emojis on page 82](#).

9.3.4 Reserved characters in the alphabet file

If you add a reserved character to the alphabet and was able to bring it back in your results, it's because dtSearch treats reserved characters as operators regardless of what you set in the Alphabet file. Consider how those operators act when you determine whether a solution works.

For example, you added the % to your Alphabet file and removed it from the Ignore list, and you were able to bring back *apple%*.

The % is the fuzzy operator, meaning you can have any one character, or no character, in this spot and bring results back. This is very similar to how *, wildcard, or ?, wildcard for any single character, work. Remember that because % is no longer being ignored, it will be indexed and will show up as part of the term. The word *apple%* was returned

indirectly, because you matched the pattern apple + any indexed character. You cannot search for just % and bring back correct results.

9.3.5 Searching for symbols and emojis

This topic describes how to make symbols and emojis searchable in a dtSearch index and how to use regular expressions to search for certain symbols reserved as search operators, such as the % sign.

9.3.5.1 Searching for symbols

There are two ways to make a character searchable, depending on whether the hexadecimal character code is [less than 007F](#) or [greater than 0080](#).

Also refer to:

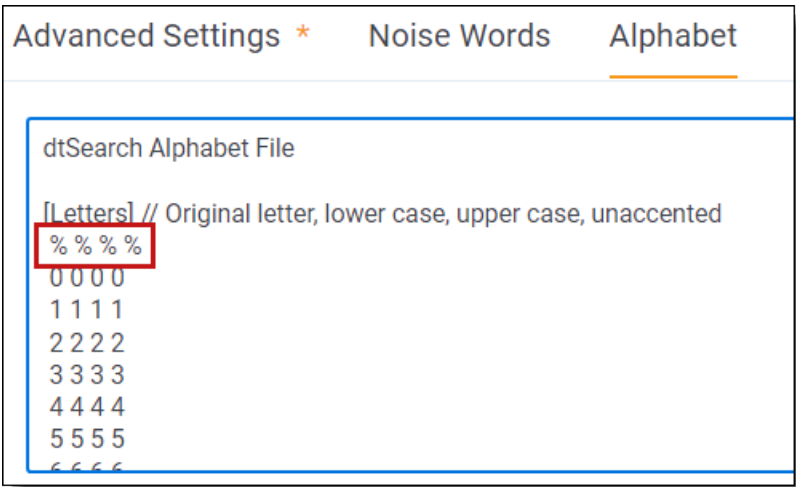
- Symbols quick reference guide
- [dtSearch Unicode values for Special Characters](#) article on Community

For hexadecimal characters less than 007F

To search for symbols with hexadecimal characters less than 007F, perform the following steps:

1. Create a dtSearch index.
2. Update the Alphabet file to include the % sign as an indexed character. For more information on the Alphabet list, see [Alphabet list on page 73](#).
 - a. Enter the following under [Letters] // Original letter, lower case, upper case, unaccented:
[Space] [%] [Space] [%] [Space] [%] [Space] [%]

Note: You must have a leading space. You cannot have a trailing space.



- b. Delete the % sign from the [Ignore] section. Take care **not** to remove the gray boxes, which represent non-printable characters. Removing them may break the index.

Before removing the %:

Note: For most symbols, once you have indexed the character in the Alphabet file, you can type them directly into the dtSearch box without using regular expressions. For example, to search for 30!, enter 30! into the dtSearch box after you have made the exclamation point (!) an indexed character. You only need to use RegEx for certain symbols.

Another example would be searching for documents that contain the term 75 %, you would enter the following in your search box, ensuring you select the proper dtSearch index.

```
"##75" w/1 "##\u0025"
```

- ## signals to Relativity to treat the string as a regular expression.
- w/1 indicates that the number be within the % sign.
- \u indicates a search using a character's Unicode value, to follow in the next four characters.
- 0025 indicates the hexadecimal Unicode value for the % sign.

For hexadecimal characters greater than 0080

For characters outside of the ASCII code range, characters greater than 0080, such as § and £, you cannot make them searchable by adding them to the [Letters] section. Instead, create an *AdditionalLetters* section at the bottom (very last line) of the Alphabet file and insert the characters' Unicode value. For example: *AdditionalLetters = 20AC 00A3 00A5*. This would make the Euro symbol (20AC), British Pound (00A3), and Japanese Yen (00A5) searchable.

The screenshot shows the 'Alphabet' tab in the dtSearch 'Advanced Settings' window. It contains three sections: '[Spaces]' with a regular expression, '[Ignore]' with a list of Unicode ranges, and '[End]' with CJK ranges. The 'AdditionalLetters' field at the bottom is highlighted with a red box and contains the text 'AdditionalLetters = 00A3 00A7'.

Note: Be sure to remove the "\u" before the four character value in the Unicode Entity field when adding to your Alphabet file.

Commonly requested searchable symbols

Below is a list of Unicode hex values necessary to make some commonly requested symbols searchable:

- **00A3** — £ (British Pound)
- **00A5** — ¥ (Yen half-width)
- **FFE5** — ¥ (Yen full-width)
- **00A7** — § (Section break)
- **00A9** — © (Copyright)
- **00AE** — ® (Registered Trademark)

- **00B6** — ¶ (Paragraph marker)
- **20AC** — € (Euro)

9.3.5.2 Searching for emojis


Relativity supports searching for emojis with dtSearch indexes.


Note the following about emojis:

- Emoji search works with file formats that store emojis as Unicode characters, such as Word documents and emails.
 - File formats that store emojis in other formats, such as storing the emoji name, are currently not supported.
 - Some examples include messages from Slack and Teams.

To run emoji searches:

1. Enter an emoji in the **search box**.

Note: You can copy and paste an existing emoji, or use the appropriate shortcut keys for your operating system. For example, in Windows, click the Windows Key  + period (.) to open the emoji sampler.



2. Click **Enter**.

Your search results appear in the document list.

Additional emojis

If you need to search for emojis specific to your organization, you can add them to the bottom of your Alphabet file using "AdditionalLetters32." After adding them to your Alphabet file, you must run a full index build.

Note: If you currently have codes listed for AdditionalLetters32, add any new codes to the end of the existing listing. For example, if you currently have "AdditionalLetters32 = 231A 231B" and want to add "23EA" and "23EB," then add them after 231B like this: "AdditionalLetters32 = 231A 231B 23EA 23EB". Do not add another "AdditionalLetters 32" line.

In the Alphabet file, add the unicode for those emojis.

To get started, use the following unicode list, which covers a large variety of emojis you may encounter:

AdditionalLetters32 = 231A 231B 23E9 23EA 23EB 23EC 23F0 23F3 25FD 25FE 2614 2615 2648 2649 264A 264B 264C 264D 264E 264F 2650 2651 2652 2653 267F 2693 26A1 26AA 26AB 26BD 26BE 26C4 26C5 26CE 26D4 26EA 26F2 26F3 26F5 26FA 26FD 2705 270A 270B 2728 274C 274E 2753 2754 2755 2757 2795 2796 2797 27B0 27BF 2B1B 2B1C 2B50 2B55 0023 FE0F 20E3 00A9 00AE 203C 2049 2122 2194 2195 2196 2197 2198 2199 21A9 21AA 2328 23CF 23ED 23EE 23EF 23F1 23F2 23F8 23F9 23FA 24C2 25AA 25AB 25B6 25C0 25FB 25FC 2600 2601 2602 2603 2604 260E 2611 2618 261D 2620 2622 2623 2626 262A 262E 262F 2638 2639 263A 2640 2642 265F 2660 2663 2665 2666 2668 267B 267E 2692 2694 2695 2696 2697 2699 269B 269C 26A0 26A7 26B0 26B1 26C8 26CF 26D1 26D3 26E9 26F0 26F1 26F4 26F7 26F8 26F9 200D 2702 2708 2709 270C 270D 270F 2712 2714 2716 271D 2721 2733 2734 2744 2747 2763 2764 27A1 2934 2935 2B05 2B06 2B07 3030 3297 3299

For additional Unicode values, use the following Unicode converter: <https://www.unicode.org/charts/>

Note: Be sure to remove the "U" before the four character value of the Unicode when adding it to the Alphabet file.

9.3.6 Setting up CJK document workspaces in Relativity

When working with documents that contain Asian character sets, specific settings help expedite your setup time in Relativity. This information can help you build a clean workspace with more accurate fields to support your Asian language document sets.

This document highlights items to be aware of as you work with workspaces that contain CJK character sets.

Note: New workspaces created in RelativityOne have extracted text automatically stored in Data Grid. Workspaces restored into RelativityOne using the ARM application will automatically have the extracted text migrated to Data Grid. To search extracted text in workspaces, you must use dtSearch or Analytics searching; you cannot use keyword search.

9.3.6.1 Requirements

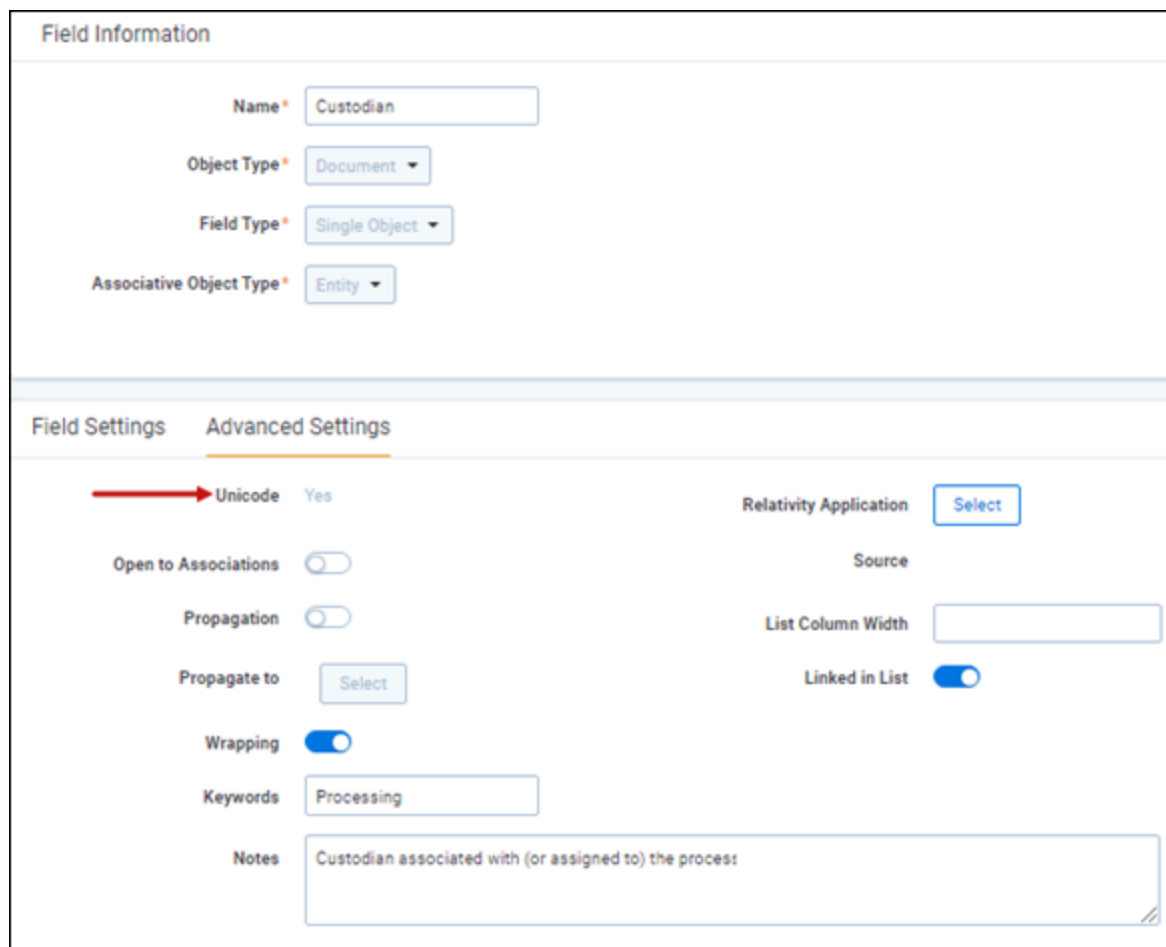
When setting up a workspace for CJK documents, there are the following requirements:

- Applicable to all versions of Relativity.
- Workspace access with the following permissions:
 - **Field:** Edit, Delete
 - **Searches:** Edit, Delete
- Import/Export application installed.

9.3.6.2 Import data with CJK characters

To import data with CJK characters:

1. On your extracted text fields, confirm the **Unicode** field is set to **Yes**.
The Unicode field is set to Yes by default.



Field Information

Name* Custodian

Object Type* Document

Field Type* Single Object

Associative Object Type* Entity

Field Settings **Advanced Settings**

Unicode Yes (indicated by a red arrow)

Open to Associations ☐

Propagation ☐

Propagate to Select

Wrapping ☒

Keywords Processing

Notes Custodian associated with (or assigned to) the process

Relativity Application Select

Source

List Column Width

Linked in List ☒

2. Set up fields such as Custodian, email fields, File Name, and Extracted Text as Unicode-compliant to capture CJK characters.
3. Open the data files that are Unicode/CJK character-compliant.
4. Use Import/Export to Import the data in a Unicode-compliant format. For more information on importing extracted text files, see Import/Export documentation on image and extracted text files.
Be sure to include any extracted text.

9.3.6.3 Searching considerations

Consider the following with SQL, dtSearch, and CJK documents. Essentially each character is its own word, or token, so proximity searching works.

Once you configure your SQL full-text language to the correct language, you can perform keyword searching and filtering in that language. Japanese in this example.

dtSearch

dtSearch can search for words in the Chinese, Japanese, and Korean languages just as it can search for words in other languages. This is because Relativity stores these languages, or converts them to, Unicode.

However, while dtSearch can search for literal word matches, or wildcard or fuzzy matches, there are some limitations to dtSearch's support for Chinese, Japanese, and Korean text.

Index

To enable support for Chinese, Japanese, and Korean (CJK) languages in Relativity, ensure that all indexed fields are Unicode-enabled. When using the default dtSearch index configuration, no changes are required to the alphabet list.

By default, the CJKRanges setting (*0e00-0e4e 3040-30ff 4e00-9fff*) ensures that each CJK character is treated as an individual word. This allows Relativity to perform accurate word breaking and searching across CJK text.

Word/Character search

You can store or convert Chinese, Japanese, and Korean text to Unicode so that you can use dtSearch to search for words in these languages just as you search for words in other languages. However, while dtSearch can search for literal word matches, or wildcard or fuzzy matches, there are some limitations on the support in dtSearch for Chinese, Japanese, and Korean text.

Those limitations include:

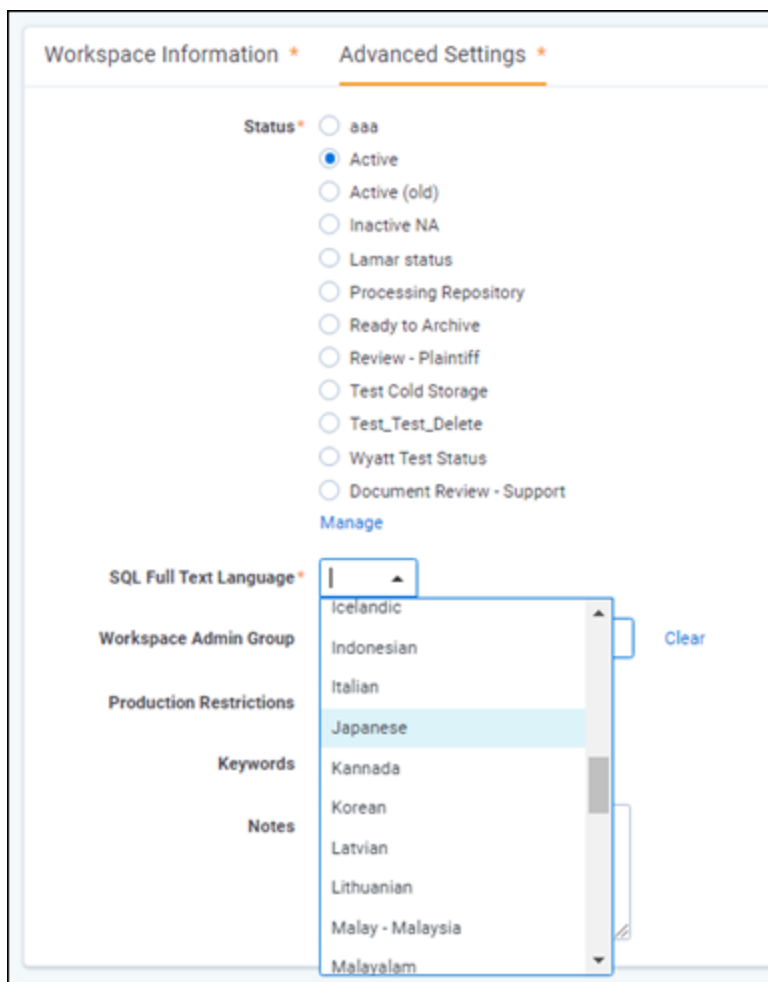
- Some documents store text in a way that does not separate the words with spaces. Instead, all of the text in a document runs together, and a language-specific dictionary is needed to find word breaks. dtSearch does not have the ability to identify word breaks in these documents because it does not include any language-specific dictionaries.
To make this type of text searchable, enable an option in dtSearch to automatically insert word breaks around Chinese, Japanese, and Korean characters. Once you enable this option, Relativity treats each character as a single word for indexing and searching purposes.
- This feature is turned on by default in dtSearch. Because each term is searchable, we recommend searching for multiple characters to assist in retrieving more accurate results. When looking for multiple characters, use the proximity connectors to assist in finding desired results.
- The same text can be presented in different ways depending on the context. dtSearch searches for a word as it is provided in the search request and does not generate additional grammatical or script variations for words in Chinese, Japanese, and Korean.
- The dtSearch engine has an API that you can use to integrate with dictionary-based language analyzers from companies such as Basis Technologies. But while the dtSearch standalone desktop does allow for integration, the instance within the Relativity environment does not support integration.

For non-Western languages, such as Chinese, Japanese, and Korean, there are additional considerations and workarounds that may provide assistance in locating search hits.

Keyword search

Within a Relativity workspace, a system admin can select the language to use in the SQL Full Text Language. SQL Full Text Language determines the correct stemming and word-break characters used in the full text index.

Note: For multiple language workspaces, Microsoft recommends setting the most complex prevalent language as the SQL Full Text language.



9.3.6.4 Language considerations

Each language has its own considerations. This section covers a few common languages and the things you should consider.

Mandarin considerations

When setting up workspaces for Mandarin, keep in the following considerations:

- Relativity treats Mandarin characters as words. Therefore the Mandarin characters do not need an extra space because each character can be compared against a list. If there is a match, then it's highlighted.
- Numbers are not treated like Mandarin characters. For example, when you use the number 54, you need everything until there is a space. Add a space to the subject so the term is 54 所 [space], then the rest of the line, it will highlight.
- The term 54 所 does not match the subject. The subject has the additional characters, hence no match.
- The Mandarin characters do not need the space because each character, such as 所, is treated uniquely, therefore, it highlights.
- **Keyword Search**—once you configure your SQL full-text language to the correct language, you can perform keyword searching and filtering in that language.
- **dtSearch**—once you set up all indexed fields to be Unicode compliant, no changes are required to the alphabet list. The default CJKRanges setting (0e00-0e4e 3040-30ff 4e00-9fff) ensures that each CJK character is treated as an individual word. This allows Relativity to perform accurate word breaking and searching across CJK text.

54 所 (other characters)	54所	No match
	所	Match
	54所 (other characters)	Match
54 所 space (other characters)	54 所	Match
	所	Match
	54 所 (other charaters)	No Match
	(other characters)	Match

Chinese considerations

When setting up workspaces for Chinese, keep in the following considerations:

- Relativity can treat Chinese characters as words. Therefore, the Chinese characters do not need an extra space because Relativity can compare each character against a list. If there is a match, then it's highlighted.
- Relativity does not treat numbers like Chinese characters. For example, when you use the number 54, followed by a Chinese character or characters, you need a match on everything until there is a space. Add a space to the search term so the term is 54 所 [space], and then the rest of the line, then it will highlight.
- The term 54 所 does not match the subject. The subject has the additional characters, hence no match.
- The Chinese characters do not need the space because each character, such as 所, is treated uniquely. Therefore, it highlights.
- **Keyword Search**—once you configure your SQL full-text language to the correct language, you can perform keyword searching and filtering in that language.
- **dtSearch**—remember that highlighting and search hits are separate implementations. Thus, there are cases where the search returns a document as a hit, but the highlighting does not show how the hit occurred.

9.3.6.5 Analytics considerations

Take the following into consideration before using Analytics:

- Use the structured data analytics language identification operation, or PLI language technique, to identify the languages of the documents in your workspace.
- Training your Analytics index in one language as opposed to multiple languages generally produces a better quality index. However, if many of your documents contain multiple languages, this can change.
- Be sure to set up your noise words in the same language as your index.
- Relativity does not provide stop words for languages other than English-configuration is left up to the administrator.

Note: Refer to the [Analytics - Support for CJK and other non-English Languages](#) Analytics - Support for CJK and other non-English Languages article for more detailed considerations.

9.3.6.6 Tokenization considerations

Languages like Chinese, Japanese, and Korean are fundamentally more difficult to process than European languages because their words are not consistently separated by spaces. Chinese and Japanese do not use spaces between

words. Korean uses some spaces, but not between every word, and inconsistently depending on the writer.

Tokenization are methods used to produce indexable units in these languages. Each resulting in different degrees of search precision. Tokenization, which is based on morphological analysis, breaks up text into real words. Morphological analysis in linguistics examines morphemes, which are the smallest units of meaning within a word , such as “un” in the word “unremarkable.”

When the system performs tokenization on the following example, there are at most three items to index. Possibly only two because the possessive marker might be treated as a noise word. A noise word in this case is the same as in English, namely a word that occurs so frequently within a language that it does not help in finding search results and is thus ignored by some search engines.

The following example breaks down how the system would tokenize the phrase *sightseeing spots in Tokyo*, which in Japanese appears as: 東京都の観光地 .

Characters	English translation
東京の観光 スポット	sightseeing spots in Tokyo
東京	Tokyo
の	(possessive marker)
観光	sightseeing

9.3.6.7 Technical workflow for tokenized character handling

When working with documents that contain certain Asian character sets, specific settings help expedite your setup time in Relativity. This information can help you build a clean workspace with more accurate fields to support your Asian language document sets.

1. In your workspace, set the Unicode field to **Yes** for your text and choice fields.
This ensures that you can store non-Western European characters in these fields.
2. Set up fields such as Custodian, email fields, File Name, and Extracted Text as Unicode-compliant to capture characters accurately.
3. Your processing vendor should have provided you with data files that are Unicode/CJK character-compliant. If this is the case, choose the encoding of the load file and any extracted text files.
Failure to set the source encoding correctly will lead to corrupt storage in your Relativity fields.

9.4 Running a dtSearch

Use dtSearch to complete stemming, fuzzy, and proximity searches on the information included in your dtSearch index.

9.4.1 Considerations

Consider the following before running a dtSearch.

- Pressing Enter between search terms in the Search Terms text box acts as spaces between terms and not as the AND or OR operator. Terms will not be searched for independently. They will be searched for as one search item.
 - Use AND or OR operators to search for terms independently.
 - For example, typing *apples*, pressing Enter, and typing *pears* searches for *apples pears* as one search item. However, typing *apples OR pears* searches for instances of the word *apples* or *pears* individually.

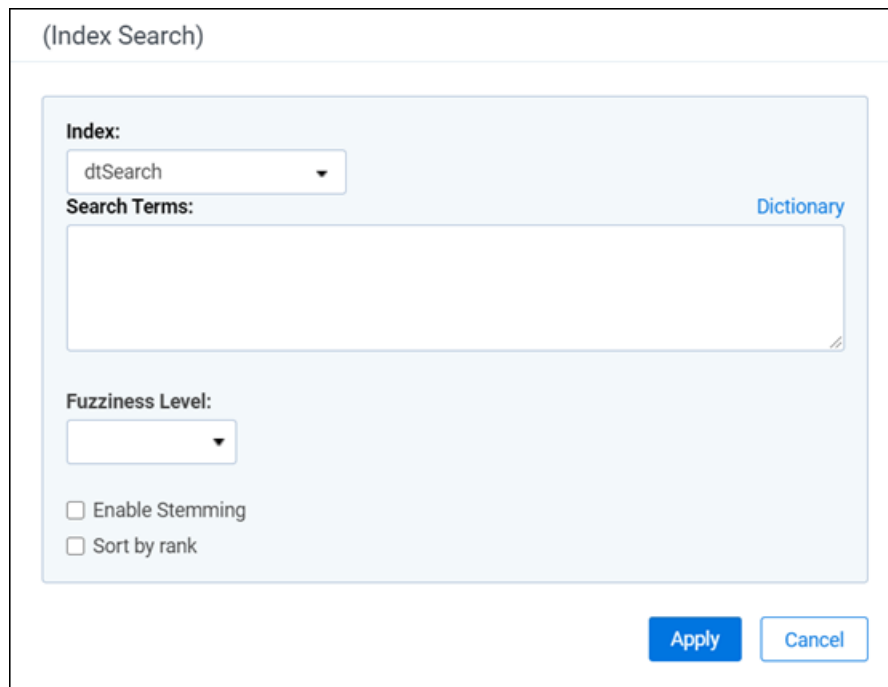
- dtSearch indexes are case insensitive by default.
 - All characters in a dtSearch index are normalized to lowercase.
 - For example, if your exact phrase search is an acronym like ACT, you must build a case-sensitive dtSearch index.
- The maximum number of search terms is 65,000 characters.
 - Words longer than 32 characters are truncated during indexing.
 - For more information, see [Searching for words longer than 32 characters on page 92](#).
- We recommend avoiding wildcards in your search terms.
 - Wildcards can slow searches.
 - Wildcards can return more documents than intended.
 - Leading wildcards should not be used. For example **ing*.
- We recommend limiting proximity operators.
 - Large numbers of proximity operators within a single query can cause slower performance. For example, *w/or PRE/*.
 - Nesting proximity operators can also cause performance problems. For example, *dog w/5 (cat w/2 bird)*,
 - Queries can be difficult to read and understand.
- We recommend running multiple queries that return small numbers of documents.
 - We recommend writing specific search terms.
 - Avoid using one broad query.

9.4.2 Running a dtSearch in the search panel

To run a dtSearch in the Search panel from the Documents list in Relativity:

1. Navigate to the search panel.
2. Click **Add Condition**.

3. Select **(Index Search)** from the Add Condition drop-down list.
The (Index Search) modal opens.




4. Select the name of your dtSearch index from the Index drop-down list.
5. Enter terms for the search in the Search Terms box. See [examples of search strings and their results](#) in the table below.

Note: Refer to [Considerations on page 88](#) for more information on search terms entries.

6. Optional fields to set:
 - **Fuzziness Level**—access fuzziness for key terms by selecting the desired level from the list.
 - **Enable Stemming**—access stemming for key terms by selecting the box.
 - **Sort by Rank**—click the box to sort the results in order of the most relevance to your search terms. The most relevant documents are listed at the top of the results set, with the least relevant documents, or the documents with the smallest number of keywords, is at the end of the list.
7. If needed, repeat the above steps to add any additional search conditions.
8. Click **Apply**.
The search terms populate in the Search builder.
9. Click **Run Search**. To stop a long running search, click **Cancel**.

9.4.3 Running a dtSearch in the Search Browser

To run a dtSearch in the Search browser:

1. Click  to access the Search browser from the Documents List.
2. Click **Create New Search**.
3. Enter required fields in the Information card.

4. Click **Add Condition**.
5. Select **(Index Search)** from the Add Condition drop-down list.
The (Index Search) modal opens.

6. Select the name of your dtSearch index from the Index drop-down list.
7. Enter terms for the search in the Search Terms box. See [examples of search strings and their results](#) in the table below.

Note: Refer to [Considerations on page 88](#) for more information on search terms entries.

8. Optional fields to set:
 - **Fuzziness Level**—access fuzziness for key terms by selecting the desired level from the list.
 - **Enable Stemming**—access stemming for key terms by selecting the box.
 - **Sort by Rank**—click the box to sort the results in order of the most relevance to your search terms. The most relevant documents are listed at the top of the results set, with the least relevant documents, or the documents with the smallest number of keywords, is at the end of the list.
9. If needed, repeat the above steps to add any additional search conditions.
10. Click **Apply**.
The search terms populate in the Search builder.
11. Click **Save** or **Save As**.
12. Select the name of the dtSearch in the Search Browser.
13. Click **Run Search**. To stop a long running search, click **Cancel**.

9.4.4 Search string examples and expected results

The following table shows search string examples and their expected results. For more information, see [Using dtSearch syntax options on page 93](#).

Search String	Returns Documents With...
apple pear	The exact phrase apple pear
"apple pear"	The exact phrase apple pear
apple AND pear	The word apple and the word pear
(apple and pear) AndAny (grape or banana)	Any document that contains apple and pear , with grape and banana also being counted as hits. Grape and banana aren't evaluated as conditions. They're simply added to the search results as optional terms to be highlighted if they appear in the document.
apple OR pear	Either apple or pear
apple W/5 pear	Apple appears within 5 words of pear
apple PRE/5 pear	Apple appears within 5 words before pear Note: Relativity does not use the POST operator. However, you can mimic this functionality by reversing the order of the terms, and using the PRE operator.
apple NOT W/5 pear	Apple does not appear within 5 words of pear
apple AND NOT pear	Apple appears but pear does not.
apple W/5 xfirst-word	Apple appears in the first 6 words of the document.
apple w/5 xlast-word	Apple appears in the last 6 words of the document.

9.4.5 Searching for words longer than 32 characters

When a word has more than 32 characters (none of which are whitespace, characters in the CJKRanges section, or symbols configured as spaces), dtSearch truncates the word and only indexes on the first 32 characters for a match. As a result, if you search for a term with 33 or more characters, dtSearch will match any document containing a word that begins with the first 32 characters of the search term. If there isn't a match for those characters, nothing will be returned.

For example, if you search for this 35-character word, "ToBeOrNotToBeThatIsTheQuestionToAsk," dtSearch will search on the first 32 characters, **ToBeOrNotToBeThatIsTheQuestionTo**, and return words that starts with those letters. The search results will include these if they are present in the searched documents:

- ToBeOrNotToBeThatIsTheQuestionTo
- ToBeOrNotToBeThatIsTheQuestionToA
- ToBeOrNotToBeThatIsTheQuestionToAs
- ToBeOrNotToBeThatIsTheQuestionToAsk

See below examples for additional references using a humorous word from the movie, "Mary Poppins": Supercalifragilisticexpialidocious.

Search	Length (in characters)	Returned results
supercalifragilisticexpialido	29	supercalifragilisticexpialido

Search	Length (in characters)	Returned results
supercalifragilisticexpialidoc	30	supercalifragilisticexpialidoc
supercalifragilisticexpialidoci	31	supercalifragilisticexpialidoci
supercalifragilisticexpialidocio	32	supercalifragilisticexpialidocio supercalifragilisticexpialidociou supercalifragilisticexpialidocious
supercalifragilisticexpialidociou	33	supercalifragilisticexpialidocio supercalifragilisticexpialidociou supercalifragilisticexpialidocious
supercalifragilisticexpialidocious	34	supercalifragilisticexpialidocio supercalifragilisticexpialidociou supercalifragilisticexpialidocious

Notes:

- If you include a wildcard (*) after the 32nd character, you will get the exact same results as if you omitted the wildcard.
- You can use the following regular expression to search your dictionary for all 32-character words: `##.{32}`

9.4.6 Using dtSearch syntax options

dtSearch includes special characters and other operators used to define search criteria. The following table lists the syntax options available for queries that run against a dtSearch index. Click the search functionality name for more details on the syntax use.

Search functionality	Special characters or operators
Auto-recognition of dates, emails, credit cards	date(), mail(), creditcard()
Boolean operators	AND, OR, NOT
Built-in search words	xfirstword, xlastword
Connector words	and, or, not, to, contains
Exact phrase - double quotes	" "
Operator precedence	()
Exact phrase - no double quotes	
Fuzzy searching	%
Noise words and the alphabet file	Noise Words, Alphabet
Numerical patterns	=
Phonic searching	#
Regular expressions (Redirects to another topic.)	"###"
Stemming	~
Wildcards	?, *

Search functionality	Special characters or operators
W/N operator	W/N
Proximity with terms order	PRE
Words and phrases	

For the list of the special characters recognized as spaces that cause word breaks, see [Alphabet file](#).

9.4.6.1 Auto-recognition

Auto-recognition provides you with the ability to search for various date formats, email addresses, and credit card numbers. However, it can dramatically affect indexing and searching performance. You must activate auto-recognition before you can use it in your workspace. Contact your system administrator for more information.

Date recognition

Date recognition searches for strings that appear to be dates. It uses English-language months, including common abbreviations, and numerical formats. For example, dtSearch recognized the following date formats:

- January 15, 2006
- 2006/01/15
- 1/15/06
- 1-15-06
- The fifteenth of January, two thousand six

Note: The short month format, *Jan*, *Feb*, and so forth, can be problematic, and is occasionally rejected by Relativity. The recommendation is to stick with the full name of the month to avoid any errors. For example *January*, *February*, and so forth.

Note the following date and date range search strings:

- To search for a date, enter a date expression between the parentheses in the string `date()`; for example, `date(january 10 2006)`.
- To search for range of dates, enter a date range between the parentheses in the string `date()`; for example, `date(january 10 2006 to january 20 2006)`.
- To search for a range of dates near the word apple, enter `date(january 10 2006 to january 20 2006) w/10 apple`.
- dtSearch does not support unterminated date ranges. To search for any date after or before a particular date, enter a bounded range with a maximal or minimal value for the bounds. The maximum value for a year is 2900, and the minimum value is 1000. For example, `date(january 10 2006 to january 1 2900)`.

dtSearch recognizes numeric strings as dates, as long as it interpretes as a valid date. This includes formats common in the US and UK, including:

- MM/DD/YY or MM-DD-YY
- MM/DD/YYYY or MM-DD-YYYY
- DD/MM/YY or DD-MM-YY
- DD/MM/YYYY or DD-MM-YYYY

In the case of ambiguous dates, such as 01/05/10, dtSearch defaults to MM/DD/YY. If the date contains words dtSearch converts the words to a numeric value to help interpret the date. For example, 30 must be a day and not a month, and 2015 must be a year, not a day or month.

Email address recognition

Email address recognition searches for text with the syntax of a valid email address, such as sales@example.com. With this feature, you can search for a specific email address regardless of the alphabet settings for "@", ".", or other punctuation in the email address.

You can also use the word listing functions in dtSearch to enumerate all email addresses in a document collection. You must include either the * or ? wildcard expression to enumerate all email addresses in a document collection.

- *mail(sales@example.com)* returns the exact email address: sales@example.com.
- *mail(sa*@example.com)* returns variations of the email address: sal@example.com; sales.sa@example.com.

Credit card number recognition

Credit card number recognition searches for any sequence of numbers that matches the syntax for a valid credit card number issued by a major company, such as Visa and MasterCard. dtSearch recognizes a credit card number regardless of the pattern of spaces or punctuation embedded in the number:

- 1234-5678-1234-5678
- 1234567812345678
- 1234 5678 1234 5678

Credit card issuers use numerical tests to exclude sequences of numbers that are not valid credit card numbers. Since these tests do not detect all invalid numbers, the feature for credit card number recognition may find additional invalid numbers.

To search for a credit card number, enter a credit card number between the parentheses in creditcard() as exemplified in creditcard(1234*).

9.4.6.2 Boolean operators

The dtSearch engine supports Boolean operators, including AND, OR, and NOT. You can use these operators to connect multiple phrases or terms in a single search expression.

Note: When using Boolean operators in a proximity search, dtSearch includes noise words. Although not searchable, a proximity searches still counts noise words.

AND operator


Note: For details on parsing proximity and Boolean strings in search conditions, see [dtSearch - How are Proximity and Boolean \(AND/OR\) parsed in search conditions?](#) knowledge base article on the Relativity Community site.

When you use the AND operator to connect expressions, only documents that contain all the expressions in the search string return in the result set. The following search strings illustrate how to use this operator:

- apple pie AND poached pear retrieves any documents that contain both phrases.
- (apple or banana) AND (pear w/5 grape) retrieves any documents that contain apple or banana AND contain pear within five words of grape.

The following graphic depicts what documents return when you use the AND operator in a dtSearch string:

AND Operator



9.4.6.3 Built-in search words

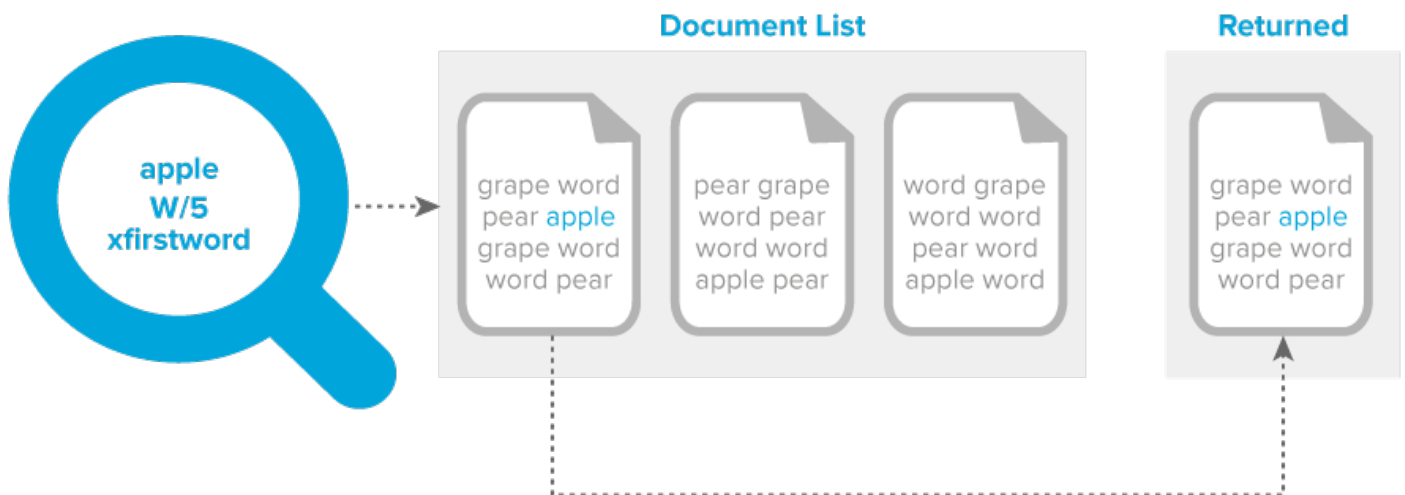
dtSearch includes the following built-in search words:

You can use these terms to limit a search to the beginning or end of a file. For example, apple W/10 xlastword searches for apple within 11 words of the end of a document.

- **xfirstword**—marks the beginning of a file.

The following graphic depicts what documents return when you use the W/# xfirstword operator in a dtSearch string:


W/# xfirstword Operator



- **xlastword**—marks the end of a file.

The following graphic depicts what documents return when you use the W/# xlastword operator in a dtSearch string:

W/# xlastword Operator



9.4.6.4 Connector words

The dtSearch connector words include:

- and
- or
- not
- to
- contains

To search for a phrase that contains one of the dtSearch connector words, quote a connector word or the phrase it is in, or put a tilde after the connector. The following search strings work in returning phrases that contain connector words:

- "clear and convincing evidence"
- not~ relevant
- "whether or not John wants to"

Note the following:

- Adding a ~ after a connector word prevents dtSearch from recognizing the word as a connector but does not otherwise affect the search. The ~ character after a word tells dtSearch to apply the stemming rules to it. Because the stemming rules included with dtSearch do not modify short words, the ~ does not change the outcome of a search for and, or, not, or to.
- The noise word list includes connector words such as *and* and *not* by default. All these words are noise words and you must remove these words from the list to make dtSearch index these files. See Creating a dtSearch index for details.
- Relativity does not support the dtSearch contains connector functionality. The word *contains* must also be treated separately to be searched.

9.4.6.5 Exact phrase - double quotes

You must use double quotes when searching for exact phrases that contain dtSearch operator reserved words, such as the Boolean connectors AND, OR. For example:

Note: Connector words such as and and not are in the noise word list by default. All these words are noise words and you must remove these words from the list to make dtSearch index these files

Search string: *clear and present danger*

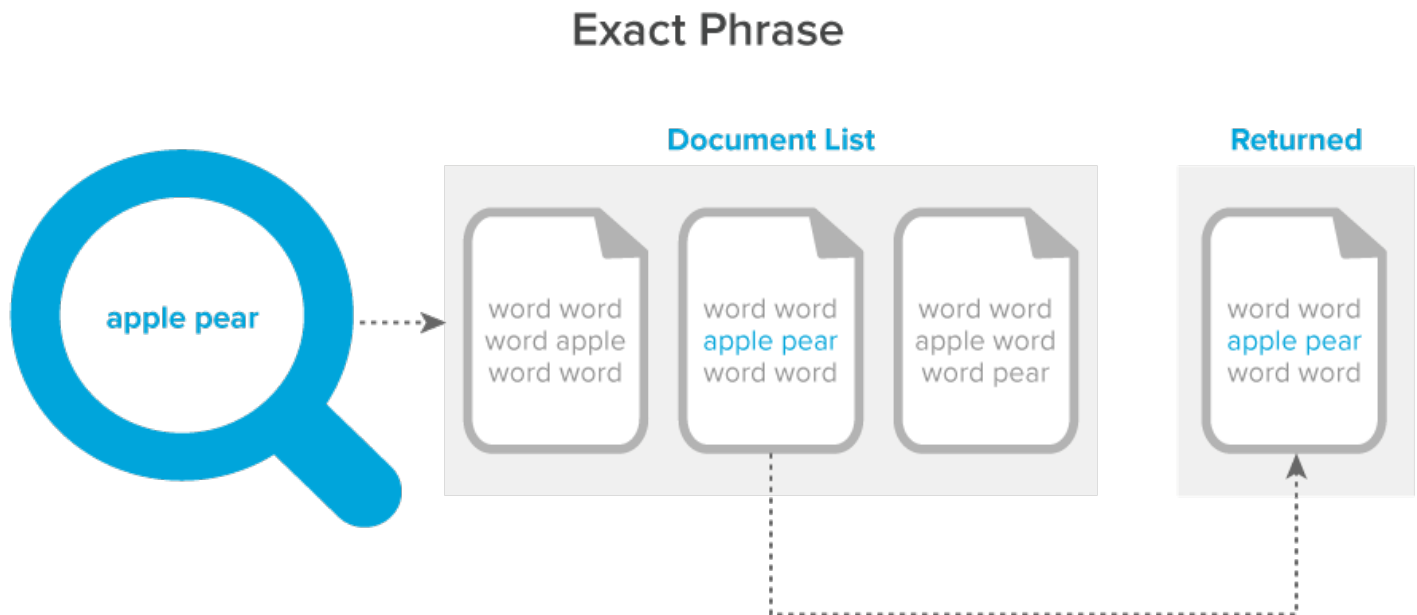
- Returns documents that contain both the word clear and the phrase present danger.
- If you need to return documents that contain the exact phrase clear and present danger, you must:
 1. Remove the word and from the dtSearch noise words list.
 2. Surround the search string with "double quotes" so that the word AND is not treated as a Boolean connector.

Search string: *"clear and present danger"*

- Returns the exact phrase clear and present danger.

Note: Do not confuse the parentheses function for order of preference with the double quotes function.

The following graphic depicts what documents return when you execute an exact phrase search:



AndAny operator

You can combine a search for required search terms with other optional terms. The words before the AndAny connector constitute required search terms, and the words after the AndAny connector are optional. A document only returns if it contains at least the required search terms. For example, (apple and pear) AndAny (grape or banana) would find any document that contains apple and pear, with grape and banana also counts as hits only if apple and pear are also present in the document.

The following example further explains the AndAny operator:

You have three documents, each containing the terms specified below:

- Document 1: Apple
- Document 2: Apple, Grape, Pear
- Document 3: Grape, Pear

Note the following behavior:

- When you search for the term *apple*, documents 1 and 2 return.
- When you search for the string *apple AND pear*, only document 2 returns.
- When you search for the string *apple AndAny pear*, documents 1 and 2 return.

The following graphic depicts what documents return when you use the AndAny operator in a dtSearch string:




OR operator

When you use the OR operator to connect expressions in a search string, documents that contain one or more of these expressions return in the result set. For example, the search string *apple pie or poached pear* returns documents that contain apple pie, poached pear, or both phrases.

The following graphic depicts what documents return when you use the OR operator in a dtSearch string:

OR Operator



NOT operator

In a dtSearch, you can use the NOT operator at the beginning of a search expression to negate its meaning and exclude documents from a result set. For example, the search expression applesauce and NOT pear returns documents that contain the word applesauce, but not those documents that contain both the words applesauce and pear.

- **NOT operator as a standalone**—you can use the NOT operator by itself at the beginning of a search expression. For example, the search expression NOT pear returns all the documents that do not contain the word pear. The search expression NOT (apple w/5 pear) returns all the documents that do not contain the word apple within five words of pear. Other examples:
 - NOT (apple or pear) returns every document that does not have apple or pear in it.
 - NOT (apple and pear) returns documents where apple and pear do not appear together in the same document. It returns all other documents including documents with the word apple and documents with the word pear. It does not return documents that include both terms.
- **NOT operator as a connector**—when the NOT operator appears in the middle of a search expression, you must also use either AND or OR. For example, the search expression apple OR NOT pear returns all the documents that contain the word apple and those that do not contain the word pear.

Note: You can also use NOT in a proximity search as illustrated by the NOT W/N, NOT Within N words, operator.

- **AND NOT operator**—you can use the AND NOT operator to develop queries for documents that include the first expression but not the second expression. For example, you may want to query for email messages that have Ryan as the author, but do not have Will as the recipient. The following record illustrates these conditions:


Document	OCR	Recipient	Author
AS00001	From: Ryan To: Will	Will	Ryan

- You can perform a dtSearch using the search string Ryan AND NOT Will and return results that do not include document AS00001.
- The dtSearch engine combines into a single pool the text for all fields identified for inclusion in an index. A search string using the AND NOT operator queries the index that includes the combine text from all

indexed fields, rather than querying the content of individual fields. This behavior ensures consistent result sets when querying with the AND NOT operator.

Note: A keyword search is an SQL full text search, which queries individual fields. Keyword searches do not return the same results as dtSearch when using the NOT operator to query across multiple fields. See [NOT operator](#).

The following graphic depicts what documents return when you use the AND NOT operator in a dtSearch string:



Operator precedence - no parentheses

The precedence, or order of evaluation, determines how a group of expressions evaluates in a query.

Note: By default, dtSearch evaluates OR expressions before AND expressions: A AND (B OR C). Unlike dtSearch, the order of precedence for a keyword search evaluates AND expressions before OR expressions: (A AND B) OR C. See Keyword search.

Evaluation order for the search string: *apple AND pear OR grape*

- pear OR grape evaluates first
- AND apple evaluates second

Documents containing the following terms return:

- pear, grape, apple
- pear, apple
- grape, apple

Operator precedence - with parentheses

Parentheses allow you to group expressions and control the order of query string execution where the query string contains both AND and OR operators. dtSearch requires both AND and OR operators for the parentheses to affect query results and ignores parentheses when the query string does not contain both operators.

For query strings containing both AND and OR operators, dtSearch evaluates OR first before AND. However, expressions contained within parentheses take precedence. If you want AND evaluated before OR, place the AND expression within parentheses.

Evaluation order for the search string: *grape OR (apple AND pear)*

- *apple AND pear* evaluated first as they reside within the parentheses
- *OR grape* evaluated second

dtSearch returns documents containing the following terms:

- apple, pear, grape
- apple, pear
- grape

Workaround for expressions containing only AND or OR operators

Use a proximity operator to separate query expressions. For example, insert a PRE proximity operator between each expression of the search string.

Evaluation of the search phrase: *(grape OR apple) PRE/1 (banana OR pear)*

dtSearch returns documents containing the following terms:

- grape banana
- grape pear
- apple banana
- apple pear

Evaluation of the search phrase: *(grape OR apple) (banana OR pear)*

dtSearch ignores the parentheses and analyzes the query as *grape OR apple banana OR pear* and returns documents with the following terms:

- grape
- apple banana
- pear
- grape, apple banana, pear

9.4.6.6 Exact phrase - no double quotes

Searching for words next to each other with no operator between them constitutes an exact phrase in dtSearch. For example, if you search for apple pear, dtSearch returns documents that contain the exact phrase apple pear. There is no rule that requires double quotes around a phrase of any number of words. You only need to use double quotes when searching for a word that is a dtSearch operator. For more details, see [Exact phrase - double quotes on page 98](#).

Search string: *pear orange*

- Returns the exact phrase: *pear orange*
- Does not return standalone word: *pear*
- Does not return standalone word: *orange*

Search string: *apple grape banana*

- Returns the exact phrase: *apple grape banana*
- Does not return partial phrase: *apple grape*
- Does not return standalone word: *grape banana*

9.4.6.7 Fuzzy searching

Using the dtSearch engine, you can perform fuzzy searches, which return documents containing spelling variations of a specified term. You may want to use fuzzy searching when querying documents that contain misspelled terms, typographical errors, or you have scanned with Optical Character Recognition (OCR).

Note: Fuzzy searching is not supported in Search Terms Reports (STR). You must use the Dictionary function to perform fuzzy searching. You would then add the desired terms to the STR. For information on how to use the Dictionary, see [Running a Dictionary search](#).

The percent sign (%) is the character used for fuzzy searches. The number of % used indicates how many characters in the search term dtSearch engine ignores when it runs the query. The position of the % indicates the number of characters from the beginning of the term that must match exactly with words in the result set. The following search strings illustrate how to use this character:

- app%ly indicates that a matching word must begin with app and differ from apply by only one character.
- a%%pply indicates that a matching word must begin with a and differ from apply by only two characters.

Using the fuzziness operator and fuzziness level option

In Relativity, you can use the fuzziness character (%) or the **Fuzziness Level** menu to perform fuzzy searches. The availability of these search options depends on the location where you are running a dtSearch:

- **Documents tab**—when you select a dtSearch in the **Search With** option, you can use the fuzziness character (%). See [Running a dtSearch](#).
- **Dictionary Search**—when you click the **Dictionary** link, you can use the fuzziness character (%) and the **Fuzziness Level** menu on the Dictionary Search dialog. See [Running a Dictionary search](#).

In the **Fuzziness Level** menu, you can select a value from 1 to 10, which applies to all terms in the text box. Larger numbers return terms with more variation. We recommend using values between 1-3 for moderate error tolerance. The following table describes the expected results for sample settings.

Fuzziness level	Description of search results
Blank	Only returns the entered term.
1	Returns slight variations of the entered term.
4	Returns multiple variations of the entered term.

- **Saved Search**—when you create a saved search, you can use the fuzziness operator (%) and the **Fuzziness Level** menu when you add a dtSearch index condition or by clicking the **Dictionary** link. The **Fuzziness Level** menu in a saved search uses the same settings as described above. See [Saved search](#).

Note: The **Fuzziness Level** menu is independent of the fuzziness (%) character that you can enter in the text box. A search for appl% without a *Fuzziness Level* setting may return documents containing apple or apply, since these terms have the stem appl and differ by one character.

Fuzzy searching uses term length and fuzziness level to decide how many % characters to add. This is not a straight level to character match. This means a level seven fuzziness search does not necessarily mean up to seven additional characters return.

9.4.6.8 Noise words and the alphabet file

The dtSearch engine references a default list of noise words and an alphabet file when it creates a new index. The dtSearch index excludes the noise words to improve query performance and prevent unnecessary index growth. When you run a query, dtSearch ignores words such as AND, THE, and WILL. The alphabet file determines how queries handle characters and spaces.

Note: If your dtSearches do not return expected results, you may want to contact your system administrator to adjust the noise word list or alphabet file.

Alphabet file

The dtSearch engine uses an alphabet file to define which characters to treat as text, cause word breaks, and ignore. System administrators can modify the default alphabet file when they create or edit a dtSearch index. See [Making a special character searchable](#).

The alphabet file determines which characters to treat as text, which cause spaces, which cause word breaks, and which to ignore. The categories of items in the alphabet file include:

- **Letters**—all searchable characters, which should include all alphabet characters, a-z and A-Z, and all digits, 0-9.
- **Hyphens**—characters removed during index creation. For example *First-Level* becomes two separate words in a dtSearch index.
- **Spaces**—characters that cause a word break. For example, the period indexes as a space character by default. Thus, dtSearch processes U.S.A. as three separate words: U, S, and A. Values listed as \## are Unicode characters. Their definitions are:
 - \09—horizontal tab
 - \0a—line feed
 - \0c—form feed
 - \0d—carriage return
 - \5c—backslash (\)

Note: Do not remove these Unicode characters from your alphabet file.

- **Ignore**—characters that dtSearch should disregard in processing text. For example, if you classify the period as ignore instead of space, then dtSearch would process U.S.A. as one word, USA.

Note: dtSearch does not recognize the underscore (_) as a space by default. Check the [Spaces] section to ensure that any character you want to treat as a word separator is properly defined in dtSearch.

Default noise word list

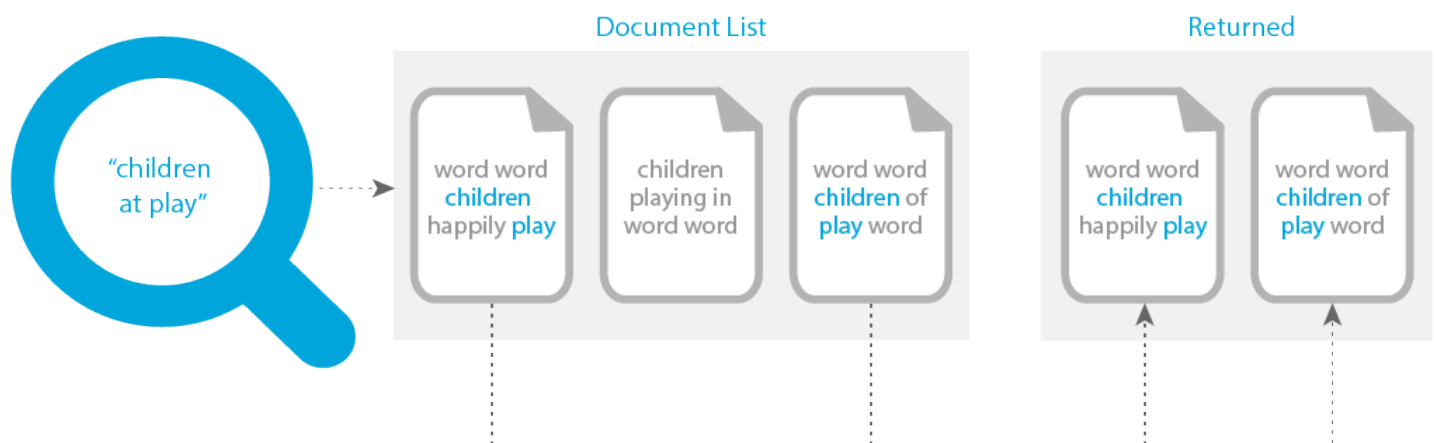
The following table shows the default noise words list. System administrators can modify this list when they create or edit a dtSearch index. Thus, if you search for a phrase that contains a term in the noise words list, you need to remove the term from the list and rebuild your index.

Begins with...	Noise words
A	a, about, after, all, also, an, and, another, any, are, as, at
B	be, because, been, before, being, between, both, but, by
C	came, can, come, could
D	did, do

Begins with...	Noise words
E	each, even
F	for, from, further, furthermore
G	get, got
H	had, has, have, he, her, here, hi, him, himself, his, how, however
I	i, if, in, indeed, into, is, it, its
J	just
L	like
M	made, many, me, might, more, moreover, most, much, must, my
N	never, not, now
O	of, on, only, or, other, our, out, over
S	said, same, see, she, should, since, some, still, such
T	take, than, that, the, their, them, then, there, therefore, these, they, this, those, through, thus, to, too
U	under, up
V	very
W	was, way, we, well, were, what, when, where, which, while, who, will, with, would
Y	you, your

The following graphic depicts what documents return when you include noise words in a dtSearch string:

Noise Words in a Phrase




Making a special character searchable

Note: You can make special characters searchable in a dtSearch index. However, you must escape some characters when using regular expressions. For more information, see the Searching for symbols recipe on the Relativity documentation site.

1. Navigate to the **dtSearch index**.
2. Click **Edit**, and then scroll down to the Alphabet section.
3. Delete the character from the current category, such as *hyphen* or *spaces*. Do not delete the category heading.
4. Enter the character you want to make searchable four times, separated by spaces under the section **[Letters] // Original letter, lower case, upper case, unaccented**.

Note: You must also begin with a space.



5. Perform a full build on the dtSearch index. The search now adds the characters you included.

Note: If you make any symbol a searchable character in your dtSearch index and then build an index on a long, uninterrupted search string, such as a file path, dtSearch truncates the string after the 32nd character. For more information, see [Searching for words longer than 32 characters](#).

9.4.6.9 Numerical patterns

To search for other numerical patterns, such as social security numbers, you can use the = wildcard, which matches any single digit. For example, if you include hyphens as spaces, then the following search request would find United States social security numbers:

=== == =====

This searching pattern can return false hits. For example, no valid social security number begins with nine. However, this is the only way to get social security numbers with spaces instead of dashes.

Note: dtSearch support notes that the === == ===== notation is higher performing than a regular expression for the same pattern, assuming you are comfortable with getting some false hits.

9.4.6.10 Phonic searching

Using the dtSearch engine, you can perform phonic searching, which returns documents containing words that sound like the word you are searching for and begins with the same letter. The pound sign (#) is the character used for phonic searches when added to the front of a word. For example, a phonic search for pear also finds pair and pare.

You can also use phonic searching in Dictionary searches.

9.4.6.11 Stemming

Using the dtSearch engine, you can perform stemming searches, which return documents containing grammatical variations of a root word. Stemming limits to English only. The tilde (~) is the character used for stemming searches when added at the end of the root word. For example, a search on apply~ returns documents containing the words apply, applying, applies, and applied. After you perform a stemming search, you can enter applied in the **Find Next** box, and then click the **Find Next** icon to locate hits or grammatical variations.

Because stemming only works with the root word, it generally does not return irregular variations of a verb. For example, a search on run~ would not return ran. The dtSearch engine only supports stemming for the English language.

Using the stemming operator and enable stemming checkbox

In Relativity, you can use the stemming character (~) or the **Enable Stemming** checkbox to perform stemming searches. The availability of these search options depends where you are running a dtSearch:

- **Documents tab**—when you select a dtSearch in the **Search With** option, you can use the stemming character (~). See [Running a dtSearch](#).
- **Dictionary Search**—when you click the **Dictionary** link, you can use the stemming character (~) and the **Enable Stemming** checkbox on the Dictionary Search dialog. See [Running a Dictionary search](#).
- **Saved Search**—when you create a saved search, you can use the stemming character (~) and the **Enable Stemming** checkbox in the Search Conditions section of the form. See [Saved search](#).

The **Enable Stemming** checkbox is independent of the stemming (~) character that you can enter in the **Search Terms** box or **Dictionary Search** text box. A search for apply~ with **Enable Stemming** checkbox unselected returns apply, applied, applies, or applying. A search for apply with **Enable Stemming** checkbox selected returns the same results.

Using fuzzy searching and stemming together

With fuzzy searching *and* stemming enabled, it checks for a fuzzy match twice, once on the original term, and once comparing the stemmed word with the stemmed word in the index. A match on either counts as a hit.

9.4.6.12 Wildcards

The dtSearch engine supports special characters that you can use as wildcards. It also supports the use of leading wildcards, or those added to the beginning of a word. The following characters represent wildcards in dtSearches:

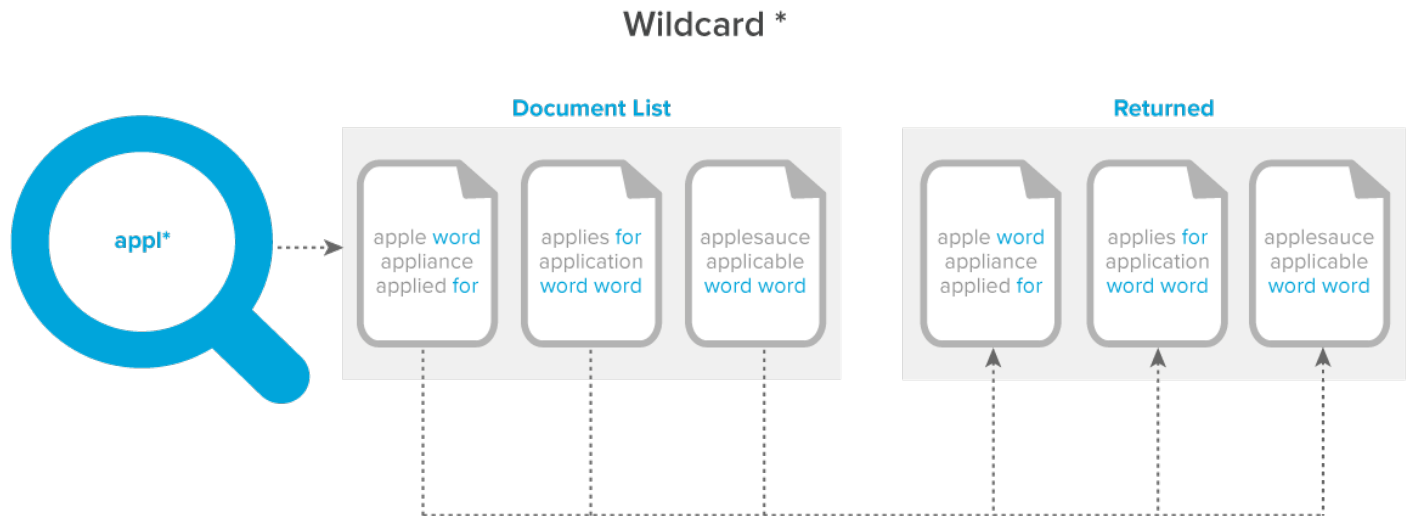
Special character	Function
?	Matches any single unknown character in the specified position in a word. Must be a character and not a blank or empty space.
*	Matches any number of unknown characters in a word, including a blank or empty space. Note: This character slows searches when used near the beginning or middle of a word.
~	Matches words containing grammatical variations of a root word. The tilde (~) is the stemming character available in dtSearches. See Stemming .
=	Matches any single numerical digit. Use multiple equal (=) signs to match a particular numerical pattern. See Numerical Patterns .

As illustrated in the following table, you can add wildcards to the root of any word to return matching terms from a dtSearch.

Sample search string	Description of search results
appl?	Matches apply and apple, but not apples.
folder?	Matches folders, but not folder.
pe?r	Matches pear, peer
?ear	Matches pear, dear, hear
appl*	Matches appl, apple, apples, application.
cipl	Matches principle, participle, disciplined.
ap*ed	Matches applied, approved, appeased

Sample search string	Description of search results
apply~	Matches apply, applied, applies.
=th	Matches 4th, 5th, 6th.
=== == =====	Matches United Stages Social Security Numbers (hyphens represented as spaces)

The following graphic depicts what documents return when you use the * wildcard in a dtSearch string:



The following graphic depicts what documents return when you use the ? wildcard in a dtSearch string:



9.4.6.13 W/N operator

You can use the W/N, within N words, operator to return documents with two words or phrases occur within a certain proximity of each other. When using Boolean operators in a proximity search with the W/N operator, dtSearch includes noise words. The N value represents the number of intervening words. For example, the search expression apple W/5 pear returns documents that contain apple only when it occurs within five words of pear. The documents returned by the search must contain the terms within the required proximity, such as five words.

The W/N operator is symmetrical. The search expression apple W/5 pear returns the exact same document as pear W/5 apple.

Note: dtSearch treats Single characters as full words when using this operator. For instance, if you search for Harry W/2 Truman, your search retrieves documents that include Harry S Truman or Harry S. Truman.

Note: Relativity does not support the WI operator. Use the W/N syntax to search for documents having words or phrases within a certain proximity of each other.

The following graphic depicts what documents return when you use the W/# operator in a dtSearch string:



NOT W/N

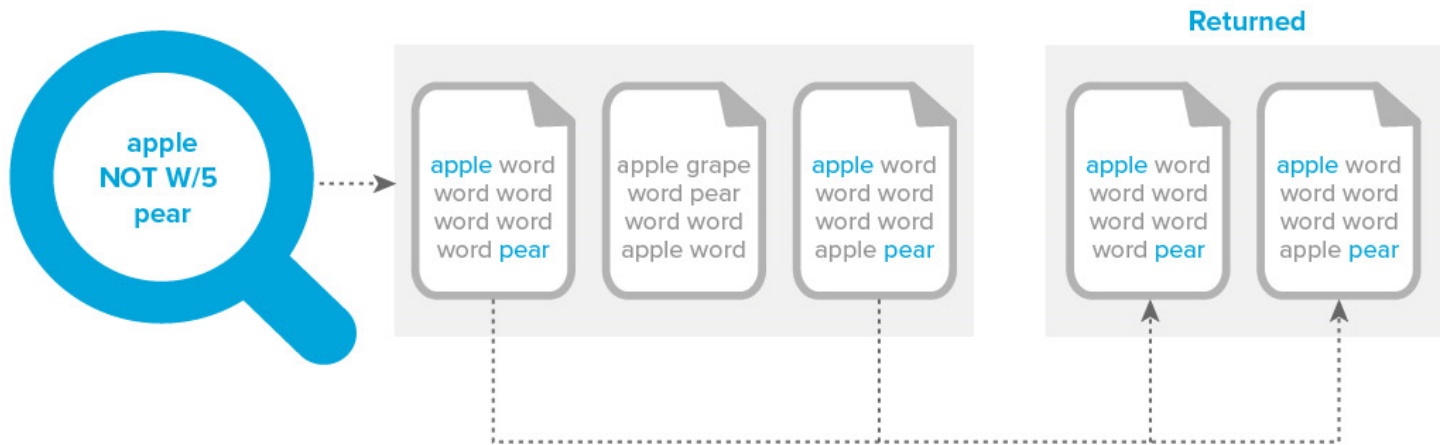
You can use the NOT W/N, not within N words, operator to exclude documents from a result set when two words or phrases are within a certain proximity of each other.

For example, the search expression apple NOT W/20 pear returns documents that contain apple when separated from pear by at least 20 words. It also returns documents that do not contain pear. Documents that contain *apple* separated from pear multiple times with varying proximity return as long as there is at least one concurrence where apple separates from pear by at least 20 words.

The NOT W/N is not symmetrical. The search expression apple NOT W/20 pear does not return the same documents as pear NOT W/20 apple.

The following graphic depicts what documents return when you use the NOT W/# operator in a dtSearch string:

NOT W/# Operator



Complex expressions

You can create complex expressions with the W/N operator by connecting words or phrases. At least one of these expressions must be a single word, phrase, or group of words and phrases connected by an OR operator as illustrated by the following:

- (apple AND banana) W/10 (pear OR grape)
- (apple AND banana) W/10 (orange tree)

Note: You can break up complex expressions with OR connectors into separate searches. Search apple w/10 "orange tree" OR banana w/10 "orange tree" to return the same results as (apple OR banana) W/10 "orange tree".

Avoid creating complex expressions that produce ambiguous results as illustrated in the following examples:

- (apple AND banana) W/10 (pear AND grape)
- (apple w/10 banana) w/10 (pear and grape)

Note: dtSearch displays a warning message when you enter an ambiguous search request.

You can also use the Boolean operators AND and OR to connect proximity expressions as illustrated in the following examples:

- (apple w/10 banana) AND (pear w/5 grape)
- (apple or banana) OR (pear w/5 grape)

Note: When connecting proximity expressions using Boolean operators, you must use parentheses.

Proximity with terms order

You can use the PRE operator to search for a word that appears within a certain number of words before another word.

For example, the search string apple PRE/5 pear returns documents where apple appears within five words before pear.

Note: Relativity does not use the POST operator. However, you can mimic this functionality by reversing the order of the terms, and using the PRE operator.

The following graphic depicts what documents return when you use the PRE/# operator in a dtSearch string:

PRE/# Operator



9.4.6.14 Words and phrases

With a dtSearch, you can use double quotes to search for a phrase. For example, the phrase *fruit salad* in the search string apple w/5 "fruit salad". The following list outlines how dtSearch queries on words or phrases with noise words or punctuation:

- **Phrases with Noise Words**—dtSearch skips any noise words in a phrase. For example, it skips of in the search string Statue of Liberty and retrieves any documents that contains statue an intervening word, and liberty.
- **Words with Punctuation**—punctuation treated as a space when inside a word. For example, dtSearch treats the search term can't as two words, can and t.
- **Numbers and Characters in Parenthesis**—you may see unexpected results when you use numbers or characters in parenthesis in a dtSearch. For example, the search term 1843 (c)(8)(ii) returns as four words.

9.4.6.15 Other considerations

- dtSearch does not recognize an underscore (_) by default. You must add the underscore to the Spaces section if you want it to be considered a space in dtSearch. For more information, see [Alphabet file on page 104](#).
- Relativity does not use the colon (:) or ampersand (&), even though considered a syntax term by the dtSearch index. If you want to use these symbols, the following applies:
 - To include either symbol in the dtSearch index, you must add it to the alphabet file.
 - To search for either symbol within file content, you must use a regular expression to define the character. Searching for special characters on their own produces incorrect results.
- dtSearch indexes are case insensitive by default. All characters in a dtSearch index normalize to lowercase. For example, if your exact phrase search is an acronym like ACT, you must build a case-sensitive dtSearch index.
- For more information on making noise words and alphabet lists searchable, see [Making the dtSearch noise word and alphabet list searchable](#).

9.5 Running a dictionary search

When you run a dtSearch, you can use **Dictionary Search** to query the index for a term and find the total occurrences and number of documents in which it occurs. The Dictionary Search dialog displays the results, including the following columns:

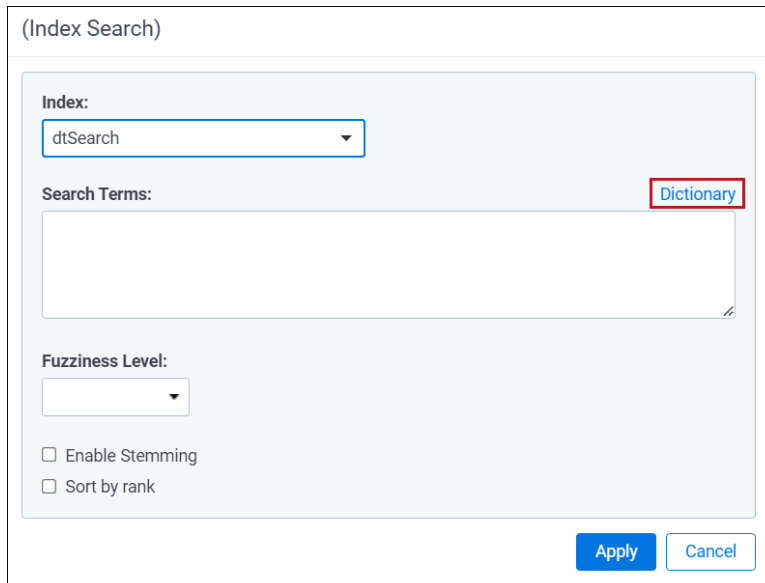
- **Keyword**—the word contained in the index. If you use fuzziness or stemming in your search, this column displays any variations of the term.
- **Total Words**—the number of times the word appears in the dtSearch index.
- **Total Documents**—the number of documents in the index that contain the word.

A dictionary search returns only the first 2,000 items in the result set. If your search returns more than that, a message displays to indicate that only 2,000 items were returned.

9.5.1 Running a dictionary search in the search panel

To run a dictionary search in the search panel:

1. Navigate to the search panel.
2. Click the **+ Condition** button.
3. Select **(Index Search)** from the Add Condition field list.
The (Index Search) modal opens.
4. Select the name of your dtSearch index from the **Index** drop-down menu.
5. Click **Dictionary**.



The Dictionary Search modal opens.

6. Enter a search term in the text box. You can enter variations of a search term as well as the wildcard (*), stemming (~), or fuzzy searching (%) operators.

Dictionary Search

Fuzziness Level

5

Enable Stemming
☒

7. (Optional) Set one of the following search operators:
 - **Fuzziness Level**—Select a value from one to ten to set the degree of variation in the terms returned. A larger number returns terms with more variation. This option is independent of the fuzziness (%) character that you can enter in the textbox on the Dictionary Search dialog. See [Fuzzy searching on page 103](#).
 - **Enable Stemming**—Select to return grammatical variations of a word. For example, a search on "apply" returns "applying", "applies" and "applied." This option is independent of the stemming (~) character that you can enter in the textbox on the Dictionary Search dialog. See [Stemming on page 106](#).
8. Click the **Search icon** to display a list of keywords and the associated totals.

Dictionary Search

Fuzziness Level

5

Enable Stemming
☒

Total items: 8

<input type="checkbox"/> Keyword	Total Words	Total Documents
<input type="text" value="Filter"/>	= <input type="text" value="Filter"/>	= <input type="text" value="Filter"/>
<input type="checkbox"/> coincidental	1	1
<input type="checkbox"/> conferentiae	1	1
<input type="checkbox"/> confidential	1	1
<input type="checkbox"/> confident	14	11
<input type="checkbox"/> confidential	193	38
<input type="checkbox"/> confidentiality	64	29
<input type="checkbox"/> confidently	1	1

9. Perform any of the following tasks with your search results:
 - **Copy Selected Keywords**—Select the checkboxes for the terms that you want to copy, and click **Copy Selected Keywords**.
The copy text modal opens.


- Use Ctrl+C to copy the terms, then close the modal.
- Click **Cancel** to close the Dictionary Search modal.
- On the (Index Search) modal, click inside the **Search Terms** box and use Ctrl+V to paste your terms.

The terms are added to the text box separated by the OR operator.

- **Sort**—Click a column header to sort in descending or ascending order.
- **Reset Columns Sizes**—Displays columns with their default widths.
- **Export**—Click to download an Excel spreadsheet containing the keywords and totals from the dictionary search.
- **Show Filters/Clear All**—See [Filters on page 11](#).


10. Click **Apply** to save your search settings.

11. Click **Run Search** to run the search against your document list.

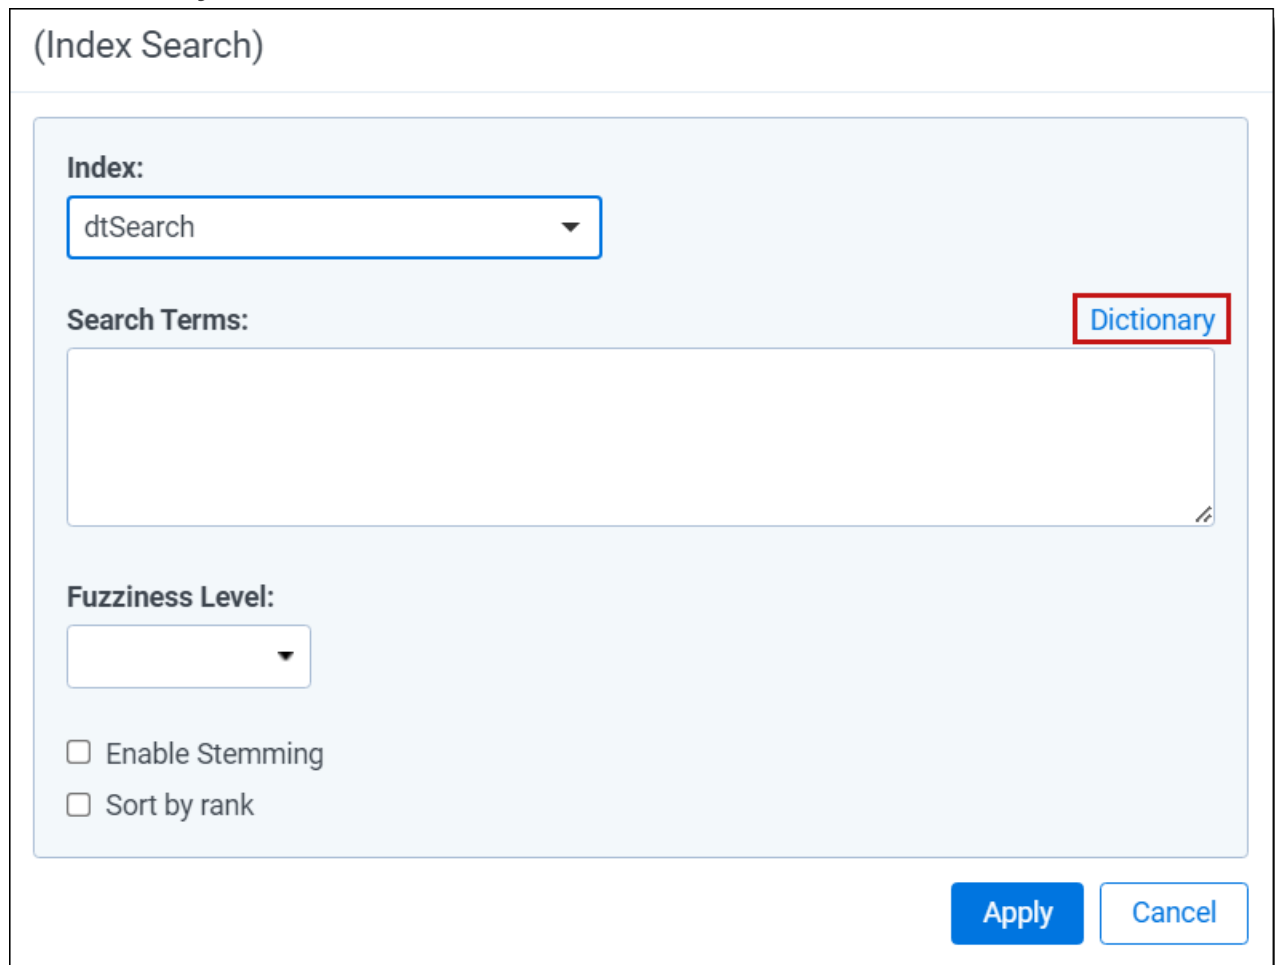


9.5.2 Running a dictionary search in the search browser

To run a dictionary search in the Search browser:

1. Click the **Saved Searches** icon () to access the Search browser from the Document list.
2. Click **Create New Search**.
3. From the **Information** tab, enter a **Name** and **Owner**.
4. Click the **+ Condition** button on the Conditions tab.
5. Select **(Index Search)** from the Add Condition field list.
The (Index Search) modal opens.
6. Select the name of your dtSearch index from the **Index** drop-down menu.

7. Click **Dictionary**.



(Index Search)

Index:
dtSearch ▼

Search Terms: Dictionary

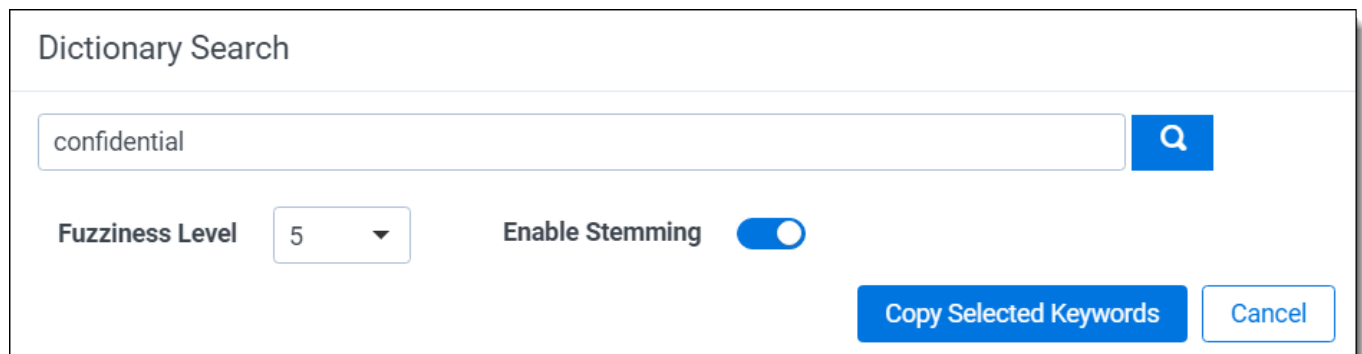
Fuzziness Level:
▼

☐ Enable Stemming
☐ Sort by rank

Apply Cancel

The Dictionary Search modal opens.

8. Enter a search term in the text box. You can enter variations of a search term as well as the wildcard (*), stemming (~), or fuzzy searching (%) operators.



Dictionary Search

confidential Q

Fuzziness Level 5 ▼ Enable Stemming ON


Copy Selected Keywords Cancel

9. (Optional) Set one of the following search operators:
- **Fuzziness Level**—Select a value from one to ten to set the degree of variation in the terms returned. A larger number returns terms with more variation. This option is independent of the fuzziness (%) character that you can enter in the textbox on the Dictionary Search dialog. See [Fuzzy searching on page 103](#).






- **Enable Stemming**—Select to return grammatical variations of a word. For example, a search on "apply" returns "applying", "applies" and "applied." This option is independent of the stemming (~) character that you can enter in the textbox on the Dictionary Search dialog. See [Stemming on page 106](#).

10. Click the **Search icon** to display a list of keywords and the associated totals.

Dictionary Search

confidential 

Fuzziness Level 5 Enable Stemming ☒

Total items: 8     

<input type="checkbox"/> Keyword	Total Words	Total Documents
<input type="checkbox"/> Filter	= <input type="text" value="Filter"/>	= <input type="text" value="Filter"/>
<input type="checkbox"/> coincidental	1	1
<input type="checkbox"/> conferentiae	1	1
<input type="checkbox"/> confidencial	1	1
<input type="checkbox"/> confident	14	11
<input type="checkbox"/> confidential	193	38
<input type="checkbox"/> confidentiality	64	29
<input type="checkbox"/> confidently	1	1

Copy Selected Keywords Cancel

11. Perform any of the following tasks with your search results:

- **Copy Selected Keywords**—Select the checkboxes for the terms that you want to copy, and click **Copy Selected Keywords**.
The copy text modal opens.
 - Use Ctrl+C to copy the terms, then close the modal.
 - Click **Cancel** to close the Dictionary Search modal.
 - On the (Index Search) modal, click inside the **Search Terms** box and use Ctrl+V to paste your terms.
The terms are added to the text box separated by the OR operator.
- **Sort**—Click a column header to sort in descending or ascending order.
- **Reset Columns Sizes**—Displays columns with their default widths.
- **Export**—Click to download an Excel spreadsheet containing the keywords and totals from the dictionary search.
- **Show Filters/Clear All**—See [Filters on page 11](#).

12. Click **Apply** to save your search settings.

13. Click **Run Search** to run the search against your document list.

- **Workspace ID**—the ID of the workspace that contains the products set used for the job.
- **Status**—the current stage of the production job.
- **Job Type**—the type of job submitted. For example, Single, Mass, or Delete.
- **Priority**—the priority setting for the indexing job.
- **Submitted Date**—the date and time the search index was run.
- **DtSearch Index Artifact ID**—the unique ID assigned to the index.
- **Job ID**—the unique ID assigned the job.

10 Searching with regular expressions (regex)

A regular expression is a form of advanced searching that looks for specific patterns, as opposed to certain terms and phrases. With regular expressions, you can use pattern matching to search for particular strings of characters rather than constructing multiple, literal search queries.

Regular expressions uses metacharacters in conjunction with a search engine to retrieve specific patterns. Metacharacters are the building blocks of regular expressions. For example, “\d” in a regular expression is a metacharacter that represents a digit character. “d” stands for the literal character, “d.” You can use regular expressions to search for social security numbers, patent numbers, URLs, email addresses, Bates numbers, and other strings that follow a specific pattern.

There are several implementations of regular expressions. The differences in implementations usually include the way special characters are handled and how character classes are treated.

10.1 Use cases for regular expressions

Regular expressions can help you in cases where you need to find different numbers that contain the same pattern.

Take, for example, the serial numbers in the first cell below. Instead of writing three literal search strings to match each serial number, you can construct one regular expression to match the serial numbers’ pattern. This single regular expression returns any document that contains any of the three serial numbers. In the second cell, there is another serial number with a slightly different pattern. By making a few adjustments to your regular expression string, your search results return documents with the new pattern.

Text	Pattern/Regular Expression
<ul style="list-style-type: none">• XFRD-8324-ERWH-3231• GHSR-3413-KBKV-8173	<ul style="list-style-type: none">• Pattern: 4 letters-4 digits-4 letters-4 digits• Regular expression: [a-zA-Z]{4}-[0-9]{4}-[a-z]{4}-[0-9]{4}
<ul style="list-style-type: none">• ABC.001.001.0001_0001• xyz.123.123.1234_1234	<ul style="list-style-type: none">• Pattern: 3 letters.3 digits.3 digits.4 digits_4 digits• Regular expression: [a-zA-Z]{3}\.[0-9]{3}\.[0-9]{3}\.[0-9]{4}_[0-9]{4}

Note: Think of each regular expression as a phrase when you construct your search string. If you switch the order of the string you will not receive the same results.

Unless you modify your dtSearch index to be case-sensitive, you cannot use capital letters when constructing a regular expression in dtSearch.

For example, if you search for strings that begin with NLRT:

- NLRT-0381
- NLRT-6334
- NLRT-9167

Use the following regular expression in Relativity: "###nlrt-\d{4}"

10.2 Regular expression metacharacters

Metacharacters are the building blocks of regular expressions. Characters in regular expressions are understood to be either a metacharacter with a special meaning or a regular character with a literal meaning.

The following are some common regular expression metacharacters and examples of what they would match or not match in regular expression.

Metacharacter	Description	Examples
\d	Whole Number 0 - 9	\d\d\d = 327 \d\d = 81 \d = 4 ----- \d\d\d ≠ 24631 \d\d\d does not return 24631 because 24631 contains 5 digits. \d\d\d only matches for a 3-digit string.
\w	Alphanumeric Character	\w\w\w = dog \w\w\w\w = mule \w\w = to ----- \w\w\w = 467 \w\w\w\w = 4673 ----- \w\w\w ≠ boat \w\w\w does not return boat because boat contains 4 characters. ----- \w ≠ ! \w does not return the exclamation point ! because it is a non-alphanumeric character.
\W	Symbols	\W = % \W = # \W\W\W = @#% ----- \W\W\W\W ≠ dog8 \W\W\W\W does not return dog8 because d, o, g, and 8 are alphanumeric characters.

Metacharacter	Description	Examples
[a-z] [0-9]	Character set, at least one of which must be a match, but no more than one unless otherwise specified. The order of the characters does not matter.	pand[ora] = panda pand[ora] = pando ----- pand[ora] ≠ pandora pand[ora] does not bring back pandora because it is implied in pand[ora] that only 1 character in [ora] can return. Quantifiers that allow pand[ora] to match for pandora is discussed below.
(abc) (123)	Character group, matches the characters abc or 123 in that exact order.	pand(ora)= pandora pand(123)= pand123 ----- pand(oar) ≠ pandora pand(oar) does not match for pandora because it's looking for the exact phrase pandora.
	Alternation—allows for alternate matches. operates like the Boolean OR.	pand(abc 123) = pandabc OR pand123
?	Question mark matches when the character preceding ? occurs 0 or 1 time only, making the character match optional.	colou?r = colour (u is found 1 time) colou?r = color (u is found 0 times)
*	Asterisk matches when the character preceding * matches 0 or more times. Note: The asterisk (*) in regular expression is different from * in dtSearch. Regular expression * is asking to find where the character, or grouping, preceding * is found ZERO or more times. dtSearch * is asking to find where the string of characters preceding * or following * is found 1 or more times.	tre*= tree (e is found 2 times) tre* = tre (e is found 1 time) tre* = tr (e is found 0 times) ----- tre* ≠ trees tre* does not match the term trees because although "e" is found 2 times, it is followed by "s", which is not accounted for in the regular expression.
+	Plus sign matches when the character preceding + matches 1 or more times. The + sign makes the character match mandatory.	tre+ = tree (e is found 2 times) tre+ = tre (e is found 1 time) ----- tre+ ≠ tr (e is found 0 times) tre+ does not match for tr because e is found zero times in tr.
.(period)	The period matches any alphanumeric character or	ton. = tone

Metacharacter	Description	Examples
.	symbol.	ton. = ton# ton. = ton4 ----- ton. ≠ tones ton. does not match for the term tones because . by itself will only match for a single character, here, in the 4th position of the term. In tones, s is the 5th character and is not accounted for in the regular expression.
.*	Combine the metacharacters . and *, in that order .* to match for any character 0 or more times. <hr/> Note: .* in regular expression is equivalent to dtSearch wildcard * operator.	tr.* = tr tr.* = tre tr.* = tree tr.* = trees tr.* = trough tr.* = treadmill

10.2.1 Regular expression quantifiers

Regular expressions uses quantifiers to indicate the scope of a search string. You can use multiple quantifiers in your search string. The following table gives examples of the quantifiers you can use in your regular expression:

Quantifier	Description	Examples
{n}	Matches when the preceding character, or character group, occurs n times exactly.	\d{3} = 836 \d{3} = 139 \d{3} = 532 ----- pand[ora]{2} = pandar pand[ora]{2} = pandoo pand(ora){2} = pandoraora ----- pand[ora]{2} ≠ pandora pand[ora]{2} does not match for pandora because the quantifier {2} only permits for 2 letters from the character set [ora].
{n,m}	Matches when the preceding character, or character group, occurs at least n times, and at most m times.	\d{2,5} = 97430 \d{2,5} = 9743 \d{2,5} = 97 -----

Quantifier	Description	Examples
		$\backslash d\{2,5\} \neq 9$ 9 does not match because it is 1 digit, thus outside of the character range.

10.2.2 Escaping regular expression metacharacters

When using regular expressions to search for a character that is a reserved metacharacter, use the backslash `\` to escape the character so it can be recognized. The following table gives an example on how to escape a reserved metacharacter when searching.

Search For	Regular Expression	Match Results
UK phone number	<code>\+[0-9]{11}</code>	>+14528280001 +38119930978 -----If the + sign is not escaped with a backslash, regular expressions treat + as a quantifier instead of the literal plus sign character.

10.3 Using regular expressions with dtSearch

You can use regular expressions with your dtSearch index to search for more complex items such as Bates numbers, zip codes, and phone numbers. You can also use regular expressions in conjunction with proximity, stemming, and fuzzy searching in dtSearch.

10.3.1 Regular expression search strings

To activate regular expressions in dtSearch, use double pound signs (`##`) at the beginning of your search string. You can start your search from the search bar on the List page, or by adding a condition from the search panel. For details on how to run a dtSearch, see [Running a dtSearch](#).

(Index Search)

Index:

dtSearch index

Search Terms:

##[a-z]{2}"##[0-9-]{5}(-[0-9]{4})"

☐ Sort by rank

Apply

Cancel

Relativity breaks down the regular expression syntax as follows:

##ExampleRegularExpressionText: The double pound signs (##) signals to Relativity that the string following ##, and enclosed by double quotes, should be interpreted as regular expression. Be sure to use straight double quotes (") and not curly quotes ("). Curly quotes may cause the regular expression to fail. You also want to avoid using capital letters in your regular expression because all characters in a dtSearch index are normalized to lowercase. You can use the Dictionary to help troubleshoot an individual regular expression. If your expression does not match in the Dictionary, it will not match in the index.

Note: Starting in Relativity 10.0.119.1, regular expression searches run from the Document List will highlight search hits in the Native Viewer for any returned documents. This does not apply to the Extracted Text mode of the Viewer.

Note: All regular expressions with dtSearch must begin with the ## call sign. If any table entries below do not include the call sign, be sure to add them to your search string before executing.

10.3.2 Regular expression metacharacters

Metacharacters are the building blocks of regular expressions. Characters in regular expression are understood to be either:

- a metacharacter with a special meaning, or
- a regular character with its literal meaning

10.3.2.1 View regular expression metacharacters examples

Note: dtSearch **does not** accept white space characters, even with regular expressions.

Metacharacter	Description	Example
\d	Whole number 0 - 9	\d\d\d = 327

Metacharacter	Description	Example
		<code>\d\d = 81</code> <code>\d = 4</code> <code>\d\d\d ≠ 24631</code> <code>\d\d\d</code> does not return 24631 because 24631 contains 5 digits. <code>\d\d\d</code> only matches for a 3-digit string.
<code>\w</code>	Alphanumeric character	<code>\w\w\w = dog</code> <code>\w\w\w\w = mule</code> <code>\w\w = to</code> <code>\w\w\w = 467</code> <code>\w\w\w\w = 4673</code> <code>\w\w\w ≠ boat</code> <code>\w\w\w</code> does not return boat because boat contains 4 characters. <code>\w ≠ !</code> <code>\w</code> does not return the exclamation point ! because it is a non-alphanumeric character.
<code>\W</code>	Symbols	<code>\W = %</code> <code>\W = #</code> <code>\W\W\W = @#%</code> <code>\W\W\W\W ≠ dog8</code> <code>\W\W\W\W</code> does not return dog8 because d, o, g, and 8 are alphanumeric characters.
<code>[a-z] [0-9]</code>	Character set, at least one of which must be a match, but no more than one unless otherwise specified. The order of the characters does not matter.	<code>pand[ora] = panda</code> <code>pand[ora] = pando</code> <code>pand[ora] ≠ pandora</code> <code>pand[ora]</code> does not bring back pandora because it is implied in <code>pand[ora]</code> that only 1 character in <code>[ora]</code> can return.

10.3.3 Regular expression groups

With regular expression groups you can match for groups of characters within a string. The following table provides examples of how to use groups in your regular expression. Groups are most useful when you use them in conjunction with alternation and quantifiers.

Metacharacter	Description	Example
<code>(abc)</code> <code>(123)</code>	Character group, matches the characters abc or 123 in that exact order.	<code>pand(ora) = pandora</code> <code>pand(123) = pand123</code> <code>pand(oar) ≠ pandora</code> <code>pand(oar)</code> does not match for pandora because it is looking for the exact phrase pandoar.

10.3.4 Escaping regular expression metacharacters

When using regular expression to search for a character that is a reserved metacharacter, use the backslash \ to escape the character so it can be recognized in its literal sense. The following table gives an example on how to escape a reserved metacharacter when searching.

Search for	Regular expression	Match results
International phone number (UK)	\+[0-9]{12}	+447700900954 +447700900312 If the + sign is not escaped with a backslash, regular expression treats + as a quantifier instead of the literal plus sign character.

10.3.4.1 Regular expression caveats in dtSearch

There are a few caveats to consider when using regular expressions in dtSearch. Consider the following caveats before constructing your regular expression.

- The metacharacter \s never matches a whitespace character in Relativity, because whitespace characters do not exist in a dtSearch index. Instead, spaces are word breaks in dtSearch. Unless you modify your dtSearch index to be case-sensitive, you cannot use capital letters when constructing a regular expression in dtSearch.

For example, if you search for strings that begin with NLRT:

- NLRT-0381
- NLRT-6334
- NLRT-9167

Use the following regular expression in Relativity: "###nlrt-\d{4}"

- You cannot search characters which are ignored during indexing, such as punctuation. To index a punctuation character, confirm that it is listed as a letter in your dtSearch alphabet file, and that it is not listed as an ignored, hyphen, or space character.

10.3.5 Common dtSearch regular expression examples

The following table includes examples of dtSearch regular expressions you can use to search for patterns in dtSearch.

Type	Regular Expression	Match Results
Bates numbers	"###rel[0-9]{7}" "###rel\d{7}"	REL0000331 REL3728948
Zip codes	"###[a-z]{2}" "###[0-9]{5}" "###[a-z]{2}" "###\d{5}"	IL 60606 MD 21218 ca 94115
United States Phone numbers	"###[0-9]{3}-[0-9]{4}" "###\d{3}-\d{4}" Note: You must make the hyphen (-) searchable in your index.	373-8837 463-9391 819-3814
United States Phone numbers	"###([0-9]{3}-)?[0-9]{3}-[0-9]{4}"	312-483-8372

Type	Regular Expression	Match Results
with or without area codes	Note: You must make the hyphen (-) searchable in your index.	463-9391
Serial numbers	<pre>###[a-z]{4}-[0-9]{4}-[a-z]{4}-[0-9]{4}"</pre> <pre>###[a-z]{4}-\d{4}-[a-z]{4}-\d{4}"</pre> Note: You must make the hyphen (-) searchable in your index.	XRFD-8324-ERWF-3231 GHSR-3413-KWEJ-8173 MPFS-1357-QEGT-9376
Dates	<pre>###[0-9]{1,2}/[0-9]{1,2}/[0-9]{2,4}"</pre>	10/17/2015 3/6/98 4/25/2006 12/04/87 95/94/93
Email addresses	<pre>##([\w_\.]+)@([\w_\.]+\.[\w_\.]{2,6})"</pre> Note: You must make the at symbol (@) and period (.) searchable in your index.	Joe.Smith426@example.com 743.MaryJane@example.com Brian.23.Voltaire@example.net.uk

11 Saved search

A saved search is a saved set of criteria that returns the latest documents that meet that criteria. For example, if you want to reference documents that contain the terms "confidential" and "property" and are also marked as Relevant, you can create a saved search with that criteria. However, saved searches can be much more complex.

In Relativity, you can create saved searches by defining custom queries and unique views, as well as by selecting public or private security settings, specific folders to query, and nested sort orders. You can also execute a search on the fly, save it for later use, or perform a combination of these tasks.

Since saved searches are executed in real-time, you save the search definition but not the results. Relativity executes the search each time you click on it in the Saved Searches browser and when you return to it after performing other tasks in the workspace. This functionality ensures that only data meeting the search criteria is returned in the result set. You can set the **Requires Manual Rerun** option to control this functionality.

Note: When you execute a saved search, Relativity first applies the conditions then related items, from the search criteria, then the filters from the item list. The only exception is when you have nested relational searches. For example, Search A which relies on the results of Search B. In this scenario, Relativity applies the innermost search conditions, in this case, Search B, then the family Search B relies on. It then applies the outer search conditions, Search A, then the family Search A relies on. Finally, Relativity applies the filters from the item list.

You can also use saved searches as the building blocks in other Relativity features. For example, you are required to select a saved search when you set up a Review Center queue, build a dtSearch or Analytics index, define an imaging set, and perform other tasks in Relativity.


11.1 Required security permissions


A user must have the following security permissions to view, edit, and add saved searches and saved search folders:

Tab/Permission	Permission	Description
Object Security/Search	View, Edit, Add	Users can view saved searches, edit existing saved searches, and add new saved searches.
Object Security/Search Container	Edit, Add	Users can edit existing saved search folders and add new saved search folders.
Other Settings/Browsers	Advanced & Saved Searches	Users can see and interact with the Saved Searches browser.



For more information, see the guide on Setting workspace permissions.

11.2 Navigating the saved searches browser

On the Documents tab, you can click  to view the Saved Searches browser. This browser provides you with features used to create, organize, edit, and perform other tasks with saved searches.



The Search Folder Tree displays the following options:

- **Create New Search** button—click this to display the Saved Search form. To display this form, you can also click any folder, including the root folder.
- **Public**  **or Private**—the icons display next to the name of a saved search to indicate its visibility.
- **Search text box**—enter the name of a search in this field to automatically filter the list as you type to the saved searches that you are looking for. See [Filtering the list of saved searches](#).
- **Display check boxes button**  —click this button to display check boxes in the list to the left of folders and searches. You can then perform mass operations for items that you check. Click the icon again to toggle them off. See [Performing mass operations on saved searches](#).
- **Search Right-click Menu**—highlight a search in the folder tree to display a right-click menu with the following options:
 - **Edit**—displays the Saved Search form, where you can modify the current settings for the search.
 - **Copy**—adds a duplicate of the search to the tree.
 - **Secure**—available on public searches, this option displays a security page so that you can override the security inherited from the workspace, or parent folder. See the Admin guide for more information on Setting permissions on Relativity objects.
 - **Delete**—permanently removes the search from the database.

- **Email Link**—opens an email message containing a link to the saved search. The Subject line is pre-populated with the following text: "Review - <Workspace Name> - <Search Name>." When the recipient clicks on the link, the saved search is displayed with the current result set.




Note: Relativity displays a permissions denied message if the recipient clicks the link to display the search but does not have access rights to it.

- **Copy Link to Clipboard**—copies a URL path to the search to your clipboard.

Note: Relativity displays a permissions denied message if the recipient clicks the link to display the search but does not have access rights to it.

- **Folder Right-click Menu**—Highlight a folder to display a right-click menu for managing folders. See [Organizing saved searches in folders](#).

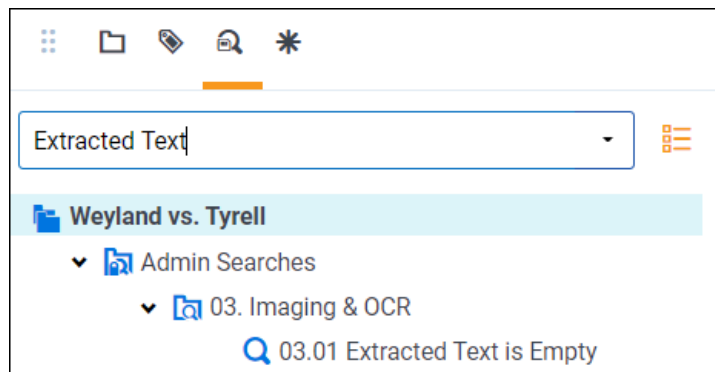
The action bar displays the following when a search is selected in the browser:

- **Show current path icon** —view the current search browser folder location for the displayed search.
- **Edit Search** —when you click this icon on the action bar, the Search Builder dialog appears where you can update search criteria.
- **Save Search** —when you click this icon on the action bar, a pop-up appears where you can select a new owner and modify the search name. See [Creating or editing a saved search](#).

11.2.1 Filtering the list of saved searches

To filter the list in the saved search browser:

1. Enter text matching the search or search folder you want to see in the Filter text box at the top of the browser. Matching searches and search folders display as you type in their respective folders.



2. Select the search you want to view.

You can also expand the Filter text box and filter on advanced searching fields, including:

- **Created By**—choose which user created the saved search.
- **Created On**—select from four different date ranges, including: Past Hour, Today, Past 7 days, and Past 30 days.
- **Last Modified By**—choose which user last modified the saved search.

- **Last Modified On**—select from four different date ranges, including: Past Hour, Today, Past 7 days, and Past 30 days.
- **Owner**—select a specific user or public.
- **Keywords**—enter keywords.
- **Notes**—enter notes.

Note: To remove your filtering from the list, delete the text that is there or click the X to the right. The list of searches will automatically update.

To access Advanced Search Filtering:

1. Navigate to the **Views** tab.
2. Locate the Advanced Search Browser View.
3. Make sure at least one of the following fields is added to the view: Created By, Created On, Last Modified By, Last Modified On, Owner, Keywords, or Notes. We recommend you add them all at once.
4. Click **Save**.


11.2.2 Performing mass operations on saved searches

Note: Add or delete permissions must be selected on the Search object for check boxes to display for the user.

The following mass operations are currently available:


- Copy
- Delete
- Move
- Export to file

To perform mass operations on saved searches in the saved search browser:

1. Click the check boxes  button to the right of the Filter text box to turn check boxes in the list of searches on.
2. Select the checkbox for an individual search or select the checkbox for a search folder to select searches inside that folder.

Note: You must expand the search folders before you can check the checkbox for the folder to allow you to examine the searches you will perform a mass operation for.

The number of selected searches appears in the drop-down list to the left of the mass operations multi-select button below. Alternatively, you can select All from the drop-down list to select all searches.



3. Once you have selected the desired searches, choose the desired mass operation, copy, delete, move, or export to file, from the mass operations multi-select button.

Note: When selecting the **move** option, sub-folders are not recreated in the destination folder.

4. A pop-up modal window appears, where you edit and confirm the details of the mass operation you are performing.
5. Click **Ok** to complete the operation.

11.3 Controlling the visibility of saved searches

On the saved search form, you can control the visibility of a search by setting the **Owner** option. New searches are private by default, making them visible only to you and Relativity administrators. In addition to owner access, users must have permissions to the **Saved Searches Browser** and at least view permissions for **Search** on the security page. For more information on workspace permissions, see the Admin Guide.

You can change the visibility of a search by selecting one of these options in the **Owner** drop-down menu:

- **Public**—makes the search available to all users with the appropriate permissions.

Note: You can configure Relativity to make your saved searches public by default. When you create a search, the **Owner** box will display Public. In **My Settings**, select **Public** in the option **Default Saved Search Owner**.

- **User Name**—select a specific user from the drop-down menu. The search will be visible only to that user and Relativity administrators.
- **Me**—click this button to reset the visibility on the search to private. Your name appears in the **Owner** box.

You define the criteria used for saved searches in the Conditions section of the Saved Search form. You can build complex queries using a combination of fields and operators that are set to required values. For information about the operators available for building these queries, as well as specific options for searching batches and developing combined searches, see [Defining criteria for saved searches](#).

11.4 Organizing saved searches in folders

You can organize saved searches by adding them to securable folders that you create and manage in the Saved Searches browser. To work with search folders, you must have the appropriate permissions for **Search Folder**, **Search**, and the **Saved Searches Browser** on the security page. For more information on workspace permissions, see the Admin Guide.

11.4.1 Adding sub-folders to the root

In the Search Folder Tree, right-click on the root folder to add sub-folders to the browser. Click **Create** to add a new folder, and name it something descriptive of its contents. To update the folder name, right-click on the folder, and click **Rename**.

11.4.2 Managing subfolders

Right-click on a folder under the root to display the following menu options:

- **Create**—adds a subfolder to the highlighted folder.
- **Rename**—makes the folder name editable. Enter new text for the name.
- **Secure**—displays a security page so that you can override the security inherited from the workspace, or parent folder. For more information on setting permissions on objects, see the Admin Guide.
- **Delete**—permanently removes all the searches and subfolders that folder contains from the database.
- **New Search**—displays the Saved Search form. See [Creating or editing a saved search](#).

11.4.3 Adding existing searches to folders

To add existing searches to a folder, left click the search and then drag and drop it into the folder. Click **OK** on the confirmation message.

Note: When you move a search, it inherits the security from the parent folder. You may want to check the security on a folder before moving a search into it.

11.5 Creating or editing a saved search

You can create and edit saved searches in the Documents tab.

11.5.1 Considerations

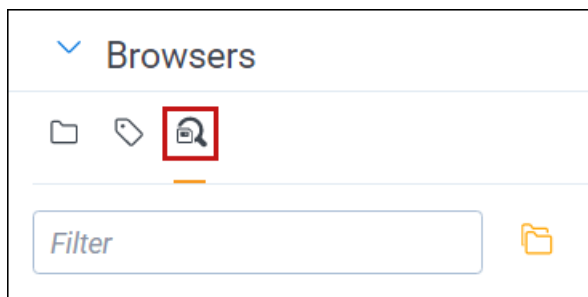
Before creating or editing a saved search, consider the following:

- We recommend avoiding nesting complex saved searches.
 - Nesting the saved search condition, can cause a timeout.
 - For example, when the saved searches used as conditions themselves contain saved search conditions.
 - If relational fields are needed, try to limit the number of nested searches or consider using a field tag.
 - Limit the use of relational fields in searches that will be used as nested saved searches later in your workflows.
- We recommend limiting the usage of the *is like* operator in search conditions.
 - It's likely to match a large number of documents.
 - Other operators return results quicker than the *is like* operator.
- If a search's execution time exceeds one minute, RelativityOne will slow it down in order to avoid impacting other users. As such, this may be a sign that your search is overly complex and would benefit from optimization.
- For any ARM restored workspaces, the Data Grid Text Migration application automatically creates a migration job to migrate the *Extracted Text* field to Data Grid. Any Extracted Text field included in a Data Grid Text Migration job that has Include in *Text Index set* to **Yes** is removed from the Keyword Search index at the beginning of the job before any text migration to Data Grid. As a result, Saved Searches using a Keyword Search and Extracted Text may fail. Be sure to update the Value and Operator fields of any Saved Searches that use the Extracted Text in a condition. For more information on text migration, see [Data Grid Text Migration](#).

11.5.2 Create a saved search

To create or edit a saved search from the Search browser, follow these steps:

1. From the Documents tab, click the **Saved Search** icon in the Browsers pane.



2. From the central pane, click **Create New Search**. To edit a search, right-click the search name, and click **Edit**.

Note: If you do not see the *Create New Search* button, you may have another saved search selected. Click the top-level folder in the browser to deselect another search.

3. Complete the fields in the **Information** section. For more details, see [Information](#).
4. From the Conditions tab, select **Condition** to add a new condition, or **Logic Group** to create a logic group. Use logic groups to group conditions. For more details, see [Conditions](#).
5. From the Fields tab, you can choose which fields you want displayed in your search results. You can also select an existing view, then add all the fields from the view in a single click. For more details, see [Search fields](#).

Note: The Default Search View controls what fields are returned, by default, on an advanced or saved search. You can always include additional fields from the advanced or saved search field selector. This view also controls the field sort order. The Default Search View is only meant to control the default fields returned when running a saved search in Relativity. Since that is its purpose in Relativity there is no way to make any changes or additions to the Conditions portion of this view.

6. Click the **Sort** tab, and then specify any sorting for the search results. For more details, see [Sort](#).
7. Click the **Other** tab to add additional information. See [Other](#).
8. Click **Save & Search**, **Search**, **Save**, or **Save As**.

Note: Selecting Save As opens a pop-up modal to update the name of the saved search.

11.5.3 Search fields

In the Saved Search window, you will want to add the following in each section.

11.5.3.1 Information

The Information fields are:

- **Name**—enter a title for the search.
- **Owner**—select an owner from the drop-down list or click **Me** to make yourself the owner.

- **Dashboard**—select an existing dashboard to link the saved search. For more information, see [Linking a dashboard to a saved search](#).
- **Requires Manual Rerun**—select this option if you want to require users to rerun a saved search when they return to it after navigating to other features in Relativity. Selecting this option only affects the search that the current user is running. It does not affect any parent or nested searches tied to the current search. If you have a search that has **Requires Manual Rerun** checked and you include it as the criteria for another search, it will rerun.

Note: The Requires Manual Rerun option is for searches that might take a long time to run, and you do not need them to run automatically when you navigate back to the saved searches. This keeps you from having to manually cancel queries before you can navigate away from that search.

- **Scope**—select one of these options to designate the document set for the search:
 - **Entire Workspace**—searches all documents within a workspace.
 - **Selected Folders**—select this option, and then click **Select Folders**. On the Select Folders pop-up, select the checkboxes for the folders that you want to search. Clear the **Include Subfolders** checkbox on the pop-up if you do not want to include subfolders. Click **Clear All** to remove all selections.
- **Notes**—any notes you want to add to give more information about the search.

11.5.3.2 Advanced

The Advanced fields are:

- **Relativity Applications**—add Relativity applications you want to associate with a saved search.
- **Keywords**—any keywords you want to give additional information about the search.
- **Query Hint**—used to optimize views. Only use the query hint if instructed by the Relativity Customer Support team. Currently, you can use Hashjoin: (*true/false*) or Maxdop: (*x*) to populate the field. You must remove query hints before using a saved search in a Relativity script.

11.5.3.3 Conditions

The Conditions fields define the criteria of the search. Click on a filter card to edit the condition or click the **x** in the top right corner to remove the condition. The equation box along the top gives you a high-level view of the conditional statement you are creating. This tab has the following controls:

- **Add Condition**—enter a field name into the Search fields text box, or select a field by scrolling through the list.

Note: If there are previously created field categories, you can select a field category from the drop-down list to conveniently filter the fields list. To learn more, visit [Field Categories](#).

- **Index Search**—select this to select a [Keyword](#), [dtSearch](#), or Analytics index, and then enter search terms to apply as a search condition.
- **Saved Search**—select this to select an existing [Keyword](#), [dtSearch](#), or Analytics saved search to apply as a search condition.
- **<field name>**—select an object field name to create a conditional expression for that field to apply to the overall search criteria.
- **Add Logic Group**—adds logic groups you can add conditions to by dragging and dropping the conditions into the logic group frames. Use the AND or OR operator to join logic groups.

- **Includes drop-down**—select an option for returning documents related to hit documents. Hit documents match the search criteria. The result set includes the related documents, but they do not need to match the search criteria. Select **No Related Items** if you do not wish to include any of these documents.
 - **Duplicates**—use this setting if you want the result set to include documents with the same MD5 Hash values as the hit documents. Relativity uses the MD5 Hash value as a unique file identifier.
 - **Family**—use this setting if you want the result set to include documents with the same group identifiers as the hit documents.
 - **<Custom Field>**—your organization may use custom related fields. Contact your system admin for additional information.

11.5.3.4 Fields

The Fields tab displays the field columns you can add or remove from your search results.

The screenshot shows the 'Fields' tab in a software interface. At the top, there are three tabs: 'Conditions', 'Fields' (which is active and has an orange star), and 'Sort'. Below the tabs, there is a search bar labeled 'Select all fields from View' with a dropdown arrow and a blue 'Add' button. The main area is divided into two panes: 'Unselected' on the left and 'Selected' on the right. The 'Unselected' pane has a search bar with 'All Fi...' and a 'Filter' button, followed by a list of fields including 'Alert', 'All Custodians', 'All Custodians (Long Text)', 'All Custodians::Custodian Deduplication Order', 'All Paths/Locations', 'All Source Locations', 'Artifact ID', 'Attachment Document IDs', 'Attachment List', and 'Attachment Name'. The 'Selected' pane has a search bar with 'Filter' and a list of fields including 'Edit', 'File Icon', 'Control Number', 'Group Identifier', 'Custodian', 'Unified Title', 'Document Extension', 'Sort Date', 'Email Subject', 'Email From', and 'Email To'. Between the two panes is a central column with four arrow buttons: a double right arrow (»), a single right arrow (>), a single left arrow (<), and a double left arrow («).

There are two fields panes:

- **Unselected**—fields you add to your search results.
To add a field, click the name, then click the **Move selected left to right** arrow in the center column. You can also click the **Move all left to the right** arrow to display all available fields.

Note: If there are previously created field categories, you can select a field category from the drop-down list to conveniently filter the fields list. To learn more, visit [Field Categories](#).

- **Selected**—fields selected that appear on the search results.
To remove a field, click the name, then click the **Move selected right to left** arrow in the center column. You can also click the **Move all right to left** to remove all fields from the search results. The position of the fields in the box is how they are ordered. You can drag and drop them to change their order.

You can also add all fields from an existing view using the **Select all fields from View** drop-down menu.

1. Click the **Select all fields from View** drop-down menu.
2. Select a view. For example, *My Checked-Out Batch*.
3. Click the **Add** button.

The fields append below any fields already displayed in the *Selected* pane. Use the move arrows to add additional fields or remove fields.



11.5.3.5 Sort

The Sort options define the default sort order used for the search results. Each row in a sort criterion has the following options:

- **Sort Field**—select a field from the left drop-down box. The search sorts on the field you select.
- **Order**—select ascending or descending from the right drop-down box.

11.5.4 Pop-up pickers

You may have the option to select values from a pop-up picker when you choose certain fields or operators in the Conditions section. For example, pop-ups are available when you select the following operators:

- Any of these
- None of these
- All of these, only for multiple object fields
- Not all of these, only for multiple object fields

See [Creating or editing a saved search](#). For information about setting batch conditions, see [Batch fields as search conditions](#).

11.5.4.1 Select items pop-up picker

Use the following general steps to select items in the picker:

1. Navigate to the Saved Search form or use the Search Conditions feature.
2. Select a **Field** option for a condition.
3. Select one or more items in the Available Items list.

A check mark indicates that you selected the item.

4. Click **Apply**.

11.5.4.2 System user fields

System user fields include the System Created By and System Last Modified By fields, which you can use in search conditions.

1. Navigate to the Save Search form or use the Search Conditions feature.
2. Select a system user field in the **Field** option for a condition.
3. Select an operator, and perform one of the following tasks:
 - If you selected **any of these** or **none of these**, select from the **Available** fields and move to the **Selected** fields column, then click **Apply**.
 - Enter the username in the text box.
4. Define any extra search criteria as needed.

11.5.4.3 Folder name field

You can select Folder Name as a field in a search condition to create more flexible queries than using the Scope section of the Saved Search form. You can combine conditions containing the Folder Name and other fields with AND or OR operators refining your search criteria.

1. Navigate to the Save Search form or use the Search Conditions feature.
2. Select **Folder Name** in the **Field** option for a condition.
3. Select an operator, and perform one of the following tasks:
 - If you selected **any of these** or **none of these**, select the checkbox next to one or more item, then **Apply**.
 - If you selected another operation, enter the folder name in the text box.
4. Define any additional search criteria as needed.

11.5.5 Search bar

The Search Bar, a UI controlled feature, exists along the top of item lists that support index search. This replicates the functionality of the Index Search condition in the Search Panel, including keyword search and dtSearch.

To run the index search,

1. Enter your search terms in the search bar.
2. Click **Search**, or hit **Enter** on your keyboard.
3. Click **Enter + Shift** to add more than one search term on separate lines.



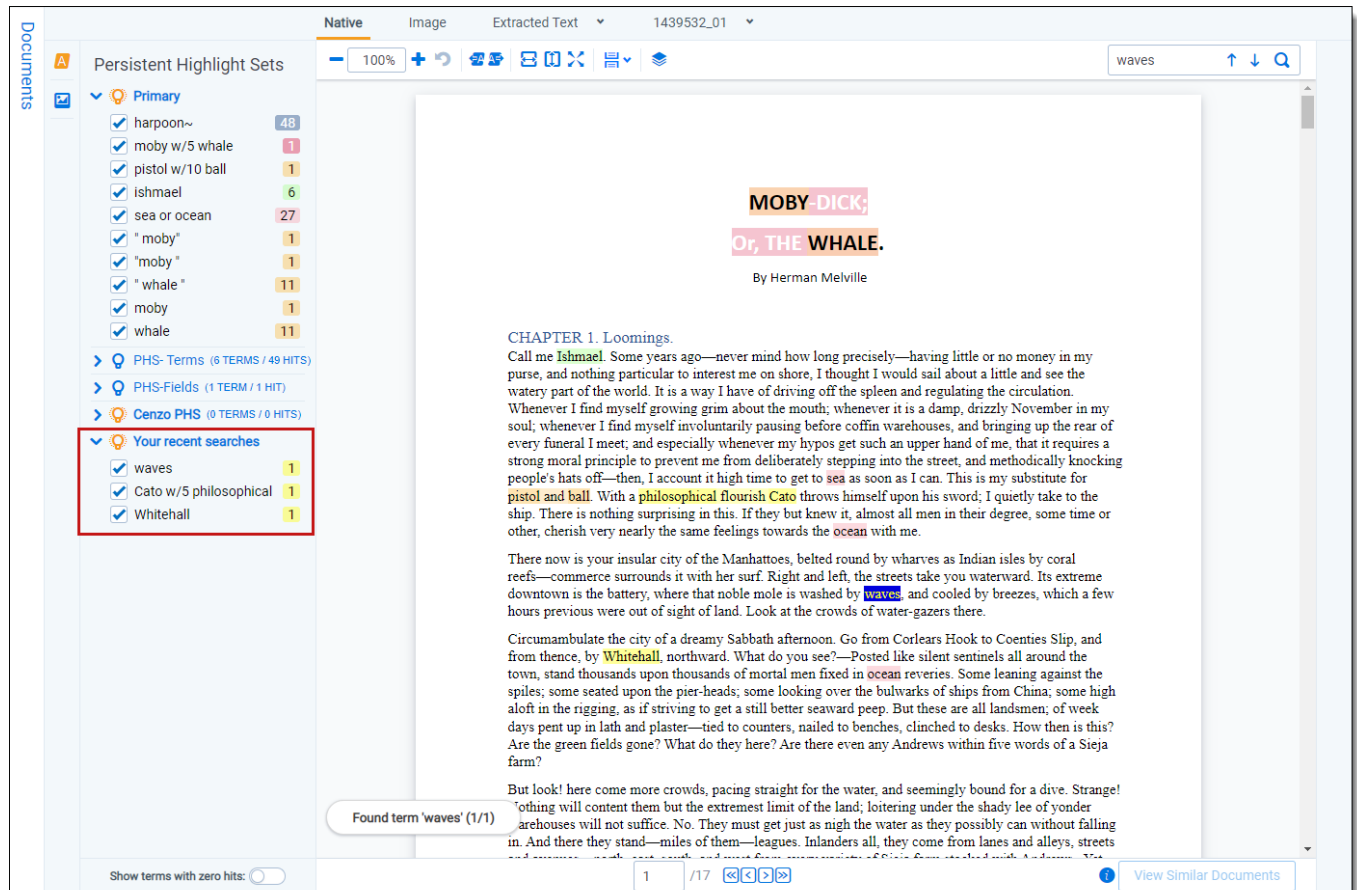
The Search Bar reflects the condition panel state after toggling index search conditions.

11.5.5.1 Recent Searches

You can use the Search Bar to generate highlights within a document in the Viewer.

To view Recent Searches within a document:

1. Navigate to **Documents**.
2. Select the desired index from the Keyword Search drop-down list.
3. Enter the desired search terms in the field to the right of the drop-list and click **Enter** on your keyboard. Keyword, proximity, fuzzy, and stemming searches can work depending on which type of search index you chose. The list displays documents that match your search criteria.
4. Click the desired document to open it in the Viewer. Highlighted terms match the entered criteria, along with any enabled Persistent Highlight Sets.



5. Optionally, click the **Show/Hide Persistent Highlight** pane to display the Recent Searches and Persistent Highlight Sets. Click the show/hide icon, light bulb icon, to toggle the Recent Searches or Persistent Highlight Sets.

Note: Recent Searches only apply when using the Search Bar or when running a saved search against a keyword index. Recent Searches do not apply when running a saved search against a dtSearch index.

11.5.6 Link a Relativity application to a saved search

To link a Relativity application to a saved search:

1. Navigate to the Search Browser.
2. Create a new search or edit an existing search.
3. Click **Add** in the **Advanced** tab to display the Select Items - Relativity Applications modal. Available applications are in the left box and Selected applications are in the right box.

4. Choose one or more applications. Use the arrows to move the applications from the available to selected boxes.
5. Click **Set**.

You can clear your options by clicking the clear button.

To unlink a Relativity Application from a saved search, click **Add** and use the arrows to remove the application(s). Please note that you cannot leave the selected applications box empty. A warning in red appears along the bottom that reads: Selection cannot be empty. Either leave at least one application linked or clear all using the clear button.

11.5.7 Link a dashboard to a saved search

To link a dashboard to a saved search:

1. Navigate to the Search Browser.
2. Create a new search or edit an existing search.
3. Choose a dashboard from the drop-down list.
4. Click **Save & Search**, **Search**, **Save**, or **Save As**.

If you switch to a different dashboard while viewing the saved search, the link breaks. There are two ways to restore this link:

- Log out of your environment, then log back in.
- Edit the search and hit **Save and Search** again.

If you delete a dashboard, Relativity removes the link in the saved search.

If you delete a dashboard with dependencies, a message appears with a list of dependencies.

11.5.8 Rerun out-of-date saved searches

You may need to rerun a saved search when you return to it after navigating to other features in Relativity. Instead of seeing your search results, you see a message indicating that your search is out of date.

Note: To enable the **Run saved search** feature, select the **Requires Manual Rerun** option in the Information section of the Saved Search form.

Perform one of these tasks:

- Click **Run saved search** to reload your search results. You can also click on the saved search in the browser to rerun the search.
- Click **Edit Search** to display the Saved Search form where you modify the search settings.

If you edit an item returned in your saved search, you need to rerun it. You must rerun the search even when the edited item still meets the search criteria, and the number of documents returned does not change.

11.5.9 Required security permissions

A user must have the following security permissions to view, edit, and add saved searches and saved search folders:

Tab/Permission	Permission	Description
Object Security/Search	View, Edit, Add	Users can view saved searches, edit existing saved searches, and add new saved searches.
Object Security/Search Container	Edit, Add	Users can edit existing saved search folders and add new saved search folders.
Other Settings/Browsers	Advanced & Saved Searches	Users can see and interact with the Saved Searches browser.

For more information, see the guide on [Setting workspace permissions](#).

11.6 Defining criteria for saved searches

You define the criteria used for saved searches in the Conditions section of the saved search form. You can build complex queries using a combination of fields and operators that are set to required values. This section provides information about the operators available for building these queries, as well as specific options for searching batches and developing combined searches. For information about building queries, see [Creating or editing a saved search on page 134](#).

11.6.1 Operators

When defining search criteria in the Conditions section of a saved search form, you use operators to determine how a field is queried for the value that you selected or entered. The operators available for a search criterion depend upon the field type:

- [Fixed-length, long, or extracted text operators below](#)
- [Whole number, decimal, and currency operators on page 144](#)
- [User operators on page 145](#)
- [Date operators on page 145](#)
- [Yes or no operators on page 146](#)
- [Single and multiple choice field operators on page 146](#)

11.6.1.1 Fixed-length, long, or extracted text operators

The following operators are available for fixed-length, long, and extracted text field types.

Operator	Returns Documents Where...
Begins with	The field begins with the entered term.
Does not begin with	The field does not begin with the entered term.
Ends with	The field ends with the entered term.
Does not end with	The field does not end with the entered term.
Is like	<p>The field contains the entered term. Blank values are filtered out, and search values can be enclosed in double quotation marks. A partial match does not require a wildcard (*).</p> <p>Relativity returns an OR operator between terms when a condition uses the "Is like" operator, and the terms are separated by a carriage return into multiple lines, or they are separated by a comma within a single line. For example, Relativity interprets a query as "Field is like Term 1 or Field is like Term 2 or</p>

Operator	Returns Documents Where...
	<p>Field is like Term 3".</p> <p>Note: For performance reasons, Relativity recommends using the "Contains" operator rather than building queries with single or multiple uses of the "Is like" operator. See FAQs for Contains and Is Like operators on the next page.</p>
Is not like	<p>The field does not contain the entered term. Wildcards (*) are already applied at the beginning and end of a term with this operator. The comments about the "Is like" operator also apply to "Is not like".</p> <p>Relativity returns an OR operator between terms when a condition uses the "Is not like" operator, and the terms are separated by a carriage return into multiple lines, or they are separated by a comma within a single line. For example, Relativity interprets a query as "Field is not like Term 1 or Field is not like Term 2 or Field is not like Term 3".</p>
Is	<p>The field value equals any of the entered items. Multiple values can be separated by a comma or carriage return. Blank values are filtered out, and search values can be enclosed in double quotation marks.</p> <p>Relativity returns an OR operator between terms when a condition uses the "Is" operator, and the terms are separated by a carriage return into multiple lines, or they are separated by a comma within a single line. For example, Relativity interprets a query as "Field is Term 1 or Field is Term 2 or Field is Term 3".</p>
Is not	<p>The field value does not equal the entered term.</p> <p>Relativity returns an OR operator between terms when a condition uses the "Is not" operator, and the terms are separated by a carriage return into multiple lines, or they are separated by a comma within a single line. For example, Relativity interprets a query as "Field is not Term 1 or Field is not Term 2 or Field is not Term 3".</p>
Is set	The field is not empty.
Is not set	The field is empty.
Is less than	The field value is less than the entered term.
Is less than or equal to	<p>The field value is less than or equal to the entered term.</p> <p>Relativity returns an OR operator between terms when a condition uses the "Is less than or equal to" operator, and the terms are separated by a carriage return into multiple lines, or they are separated by a comma within a single line. For example, Relativity interprets a query as "Field is less than or equal to Term 1 or Field is less than or equal to Term 2 or Field is less than or equal to Term 3".</p>
Is greater than	The field value is greater than the entered term.
Is greater than or equal to	<p>The field value is greater than or equal to the entered term.</p> <p>Relativity returns an OR operator between terms when a condition uses the "Is greater than or equal to" operator, and the terms are separated by a carriage return into multiple lines, or they are separated by a comma within a single line. For example, Relativity interprets a query as "Field is greater than or equal to Term 1 or Field is greater than or equal to Term 2 or Field is greater than or equal to Term 3".</p>
Contains	<p>The field includes the entered term. This operator is available for long text and fixed length text fields included in the full text index. Blank values are filtered out, and search values can be enclosed in double quotation marks. You can also use the AND/OR operators, and add wildcards (*) to the end of the search. Relativity returns an OR operator between terms when a condition uses the "Contains" operator, and the terms are separated by a comma. (You can also use the OR operator to separate search terms.) For example, Relativity interprets these queries as "Field contains Term 1 or Field contains Term 2 or Field contains Term 3", and "Field contains Term 1 or Term 2 or Term 3" respectively.</p> <p>Relativity returns an AND operator between terms when a condition uses the "Contains" operator, and</p>

Operator	Returns Documents Where...
	<p>the terms are separated by a carriage return into multiple lines, or if terms are separated by a space on a single line so "Field contains Term 1 Term 2 Term 3" is interpreted as "Field contains Term 1 AND Term 2 AND Term 3".</p> <hr/> <p>Note: The "Contains" operator works identically to keyword search, except that it searches that specific field.</p> <hr/>
Does not contain	The field does not contain the entered term. This operator is available for long text and fixed length text fields included in the full text index. Blank values are filtered out, and search values can be enclosed in double quotation marks.

FAQs for Contains and Is Like operators

You can improve your searches by understanding the differences between the "Contains" and "Is like" operators.

Why do searches using the "Is like" operator tend to run slowly?

The "Is like" operator can slow the performance of your system because it queries every document for the field specified in the condition. For performance reasons, we don't recommend building queries with single or multiple uses of the "Is like" operator. Instead, you can use the "Contains" operator, which improves performance by querying only the fields in the full-text index.

Are there any special requirements for using the "Contains" operator?

Yes, the full text-index must include the field that you want to search with the "Contains" operator.

How does query execution differ for searches with the "Contains" and "Is like" operators?

The "Contains" operator queries the SQL full text catalog, while the "Is like" operator queries the database table inside the catalog. The "Is like" operator prevents other queries from editing the table until it completes, which can negatively affect performance.

Why are different search results returned by queries using the "Contains" versus "Is like" operators?

The difference in result sets is caused by the way SQL interprets queries using these operators. An "Is like" statement appends a wildcard to the front of each query, which sometimes causes it to return more items than a query with the "Contains" operator. You must evaluate these additional items to determine if they're actually part of your expected result set or if they represent false hits for the items that you want to return.

Is there any way to enhance the performance of queries using the "Is like" operator?

Yes, you can make queries that use the "Is like" operator more efficient by creating a SQL index on the table column referenced by the query. The query can point to this index, and avoid accessing the table. For more information, contact your system admin.

How can I use the "Contains" operator to facilitate document reviews?

You can use the "Contains" operator to search fields in email headers for email addresses and other pertinent header information. To perform these searches, the Author, TO, CC, and BCC fields in email headers must be added to the full text index.

Special considerations for Data Grid-enabled fields

You can use only these operators to search Data Grid-enabled text fields:

- is set
- is not set

Note: IS SET condition operator excludes the Data Grid records where the field is null or has an empty string value.

11.6.1.2 Whole number, decimal, and currency operators

The following operators are available for whole number, decimal, and currency field types.

Operator	Returns Documents Where...
Is	The entered number is equal to the field value.
Is not	The entered number is not equal to the field value.
Is set	The field is not empty.
Is not set	The field is empty.
Is less than	The field value is less than the entered number.
Is greater than	The field value is greater than the entered number.

11.6.1.3 User operators

The following operators are available for user field types.

Operator	Returns Documents Where...
Is logged in user	The logged in user is equal to the field value.
Any of these	Any of the selected users match the field value.
None of these	The selected users do not match the field value.
Is set	The field is not empty.
Is not set	The field is empty.

11.6.1.4 Date operators

The following operators are available for date field types.

You can search on date and time but time is not displayed by default when you select a date on the calendar pop-up. For example, you can search on 10/16/2001 3:57 PM by typing in the time after your selected the date. You can also search for 2/3/10 between 4/3/10 and both 2/3/10 and 4/3/10 return in the results.

Note: When you search on a user-created date field using a relative date, Last 7 Days for example, you may see different results between a saved search and a search from the search panel if you are not in the same timezone as your Relativity instance.

Operator	Returns Documents Where...
Is	The entered date is equal to the field value.
Is not	The entered date is not equal to the field value.
Is set	The field is not empty.
Is not set	The field is empty.
Is before	The field value is before the entered date.
Is before or on	The field value is before or on the entered date.
Is after	The field value is after the entered date.
Is after or on	The field value is on or after the entered date.
Between	The field value is between the two entered dates.
Is in	The field value is within the selected range.

11.6.1.5 Yes or no operators

The following operators are available for Yes/No field types.

Operator	Returns Documents Where...
Is	The selected value (Yes or No) is equal to the field value.
Is not	The selected value (Yes or No) is not equal to the field value.
Is set	The field is not empty.
Is not set	The field is empty.

11.6.1.6 Single and multiple choice field operators

The following operators are available for single and multiple choice field types.

Operator	Returns Documents Where...
Any of these	Any of the selected choices are present in the field.
None of these	None of the selected choices is present in the field.
All of these	All of the selected choices are present in the field.
Not all of these	The selected choices are not present in the field.
Is set	The field is not empty.
Is not set	The field is empty.

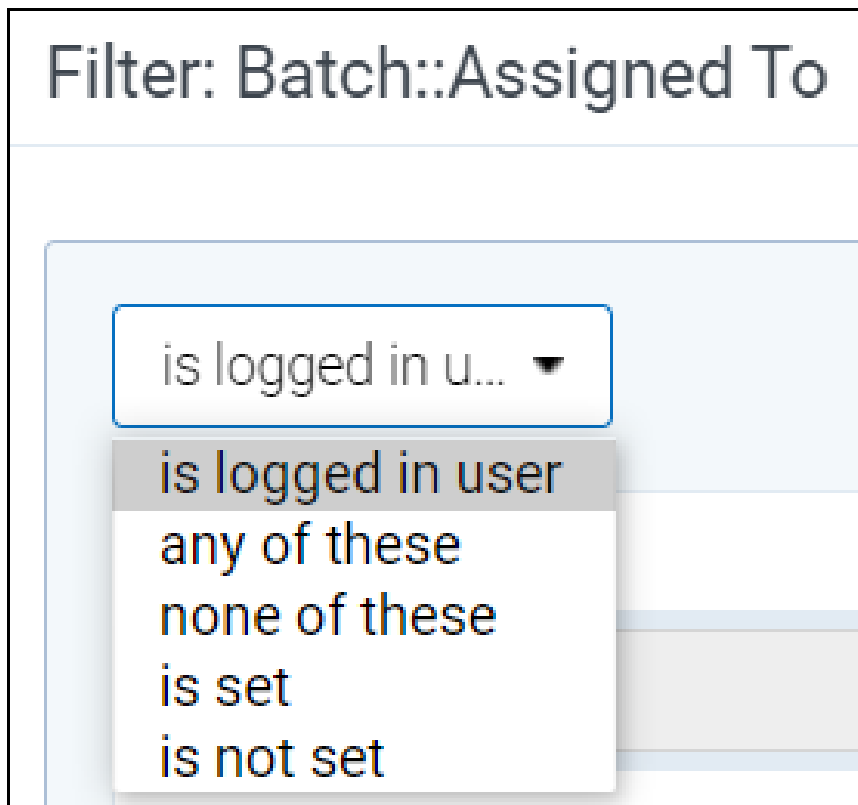
11.6.2 Batch fields as search conditions

Relativity includes several fields related to batching that you can use as conditions for searching across batch sets. To use batch fields as search conditions, follow these steps:

1. Create a new saved search. See [Creating or editing a saved search on page 134](#).
2. Select **Batch** in the Fields drop-down menu in the Conditions section of the saved search form.
3. Choose the operator you want to use.
4. Click **Add Condition** to display a modal.
5. Select one of the following batch related fields:
 - Batch
 - Batch::Batch Set
 - Batch::Assigned To

Note: You can use the *Assigned To* field to search for batches set to a specific user.

- Batch::Status



See [Operators on page 142](#) for more information.

11.6.3 Using saved searches as conditions (combined searches)

You can combine searches by selecting previously created saved searches as Field conditions. To use a saved search as a search condition, follow these steps:

1. Create a new saved search. See [Creating or editing a saved search on page 134](#).
2. Select **(Saved Search)** from the top of this list in the **Field** drop-down menu in the Conditions section of the saved search form.
3. Select an option in the **Operator** drop-down menu.
4. Click **Add Condition** to display a modal.
5. Select a saved search and click **OK**.

11.6.3.1 Preventing circular references

Relativity prevents you from creating recursive searches when you combine multiple searches as conditions in a query. For example, you might create a combined search using Saved Searches 1, 2, and 3 as follows:

- Saved Search 2 uses Saved Search 1 as a Field condition.
- Saved Search 3 uses Saved Search 2 as a Field condition.

When you edit Saved Search 1, you can't select Saved Search 2 or Saved Search 3 as Field conditions. Relativity prevents you from creating a circular reference by not listing these searches in the **Value** pop-up window.

11.6.3.2 Including related items in combined searches

You can combine saved searches to check for conflicts within related item groups. For example, a review manager may use a combined search for quality control when preparing to produce responsive documents for a case.

Use the following steps to confirm that a search for responsive documents does not include any privileged documents.

1. Create a saved search called **Responsive Check** that uses the *Includes Family* option and sets a condition on a field, such as Designation to **Responsive**.
This search specifies production criteria that return only responsive documents.
2. Create a second saved search called **Privilege Check** that uses the *Includes Family* option and sets a condition on a field, such as Privilege Description to **Privileged**, **Attorney Client**, and so on.
This search is used for evaluation purposes.
3. Create a combined search called **Conflict Check** that uses the *Includes Family* option and sets conditions for the **Responsive Check** and **Privilege Check** searches.
This quality control search determines if any privileged documents are included in the production-eligible saved search.


11.6.4 Lists as search conditions


If you've created saved lists using the Lists feature, you can add lists as criteria in a saved search.

Follow these steps to create a new saved search using a saved list of documents as the search criteria:

1. On the **Documents** tab, click the **Saved Searches browser** icon.
2. Click **Create New Search**.
3. Type a name for the saved search in the **Name** field.
4. Add a condition with the following column settings:
 - **Field**—select **Lists**.
 - **Operator**—select **these conditions**.
 - **Value**—select the following value criteria:
 - **Field**—select **Lists**.
 - **Operator**—select **any of these**.
 - **Value**—select one or more saved lists to include in the search criteria.
5. In the **Fields** category, select the fields you want to include when viewing your saved search results. Select the **Lists** field to show the list(s) with which an object is associated.
6. Click **Save** to save the search. Or, click **Save & Search** to save and execute the search.

11.7 Saving searches on the Documents tab



On the Documents tab, you can click **Save Search**  at the bottom of the screen near the mass operations. This creates a new search using the criteria that you've already set on the search panel. You can save conditional searches, keyword searches, dtSearches, or Analytics searches. Relativity also retrieves any settings that you selected for a view, sort order, or other features that control how your results appear.

Note: To use **Save Search** , you must have add permissions for **Search**, and access to the **Saved Searches Browser**. For more information on workspace conditions, see the Admin guide.


To create a saved search on the Documents tab:

1. Navigate to the Documents tab.
2. (Optional) In the **Browser** menu, select the **Folders**, **Field Tree**, or **Clusters** option.
The item list for the selected browser displays.
3. Filter on the documents in the list or select a keyword or other search option. To set search criteria, see the specific instructions for running keyword searches, dtSearches, Analytics, or filters. Your search results appear in the item list on the Documents tab.

Note: Any folder, tag, or cluster selected in the browser is included as a condition when the search is saved. The current columns, column filters, and sort order save along with any conditions set for your view (including your selection in the Related Documents drop-down menu).

4. Click .
The search builder window appears.
5. Select or enter the following required information:
 - **Name**—enter a title for the search. The title appears in the saved searches browser.
 - **Owner**—select **Public** to make the search available to all users or choose a specific user from the list. Click **Me** to select your name from the list, making the search private. (Users must have the appropriate privileges to view searches.) See [Controlling the visibility of saved searches on page 133](#).
 - **Search Folder**—click the **Ellipsis**  button to launch a pop-up window where you can save the new search to a specific folder on the saved searches browser. Highlight the folder where you want to save the search, and then click **OK** to select it.
6. Add to or modify the search criteria as needed. See Creating or editing a saved search in the new UI framework.
7. Click **Save**.

After you save your search, it appears in the saved searches browser. You can modify the search using the same

options available for saved searches. See Creating or editing a saved search in the new UI framework. Click  to display the saved search browser. To update the search, follow the same steps as those used to edit a saved search. Right-click on the search in the saved searches browser, and then click **Edit**. The saved search form pre-populates with information used in your search.

11.8 Common Saved Searches application

This topic describes how to install and use the Common Saved Searches application, which contains a group of searches that can assist you in your basic usage of Relativity.

11.8.1 Installing the application


The installation process follows the same steps used to install other Relativity applications. For more information, see Installing applications.

You must have system admin permissions to install an application. See Workspace security.

To install the application:

1. Log into the Relativity Community and search in Files for the Common Saved Searches Solutions application rap file.

Note: You must have valid Relativity Community credentials in order to download any Community file linked to the documentation site. You'll need to enter those credentials on the Community login screen if you're not already logged in. If you're already logged in to the Community at the time you click a link, the file is automatically downloaded in the bottom left corner of your screen. If you get an error message stating "URL No Longer Exists" after clicking a Community link, it may be due to a single sign-on error related to the SAML Assertion Validator, and you should contact your IT department.

2. Click **Download**.
3. Add the application to the Application Library:
 1. On the **Applications & Scripts** tab, click the **Application Library** tab.
 2. Click **Upload Application**.
 3. Click **Choose File**, navigate to and select the Common Saved Searches Application rap file, and then click **Open**.
 4. Click **Save**.
 5. Click **Install** in the Workspaces Installed section to install the application on workspaces.
6. Click  in the Workspaces field to display the Select Workspaces dialog.
7. Select a workspaces to install the application, and then click **Ok**.
8. (Optional) Click **Clear** to remove a workspace from the list.
9. Click **Save** to install the application to the selected workspaces.

Note: There may be a need to map fields if you have a template already created with the same common fields. For more information, See Mapping fields section in the Installing applications documentation.

11.8.2 Using the application

The table below lists the search name and the description of why this search is used and/or what it returns.

Search Name	Description
All Documents	A quick, easy way to return all documents in the workspace. This is useful for many reasons, including building Search Term Reports and is a method to quickly grab everything.
01.01 Processing Errors	Returns all documents in document table that have an error from processing. Use this to find items that might provide reviewers issues.
02.01 dtSearch	Returns all documents, but more importantly, returns only the extracted text field. Use to build a dtSearch index.
02.02 Structured Analytics Set	Used to check for documents with text and is used for any of the structured analytics sets.
02.03 Email Threading - Parent Emails	Search for emails, but only parent level items, not attachment emails. Previously used for email threading. Items used for threading.

Search Name	Description
02.04 Near Duplicate Analysis (No Parent Emails)	Returns all files that are identified for near duplication analysis, but are not parent emails.
02.05 Language ID or Repeated Content	Records with text to be used for Language ID or Repeated content similar to Structured set.
02.06 Analytics Not Included	Returns items not included in Analytics because of text size.
02.07 Analytics Index	Returns documents containing text and returns only extracted text for Analytics Indexes.
02.08 All docs Email Metadata fields	Returns all emails with all metadata fields returned.
03.01 Extracted Text is Empty	Includes documents not returned for use in Analytics Index for lack of text.
04.01 Parent Level Documents	Returns top-level documents that are not attachments, but can be emails or other.
04.02 Document Level Documents	Returns child-level documents.
04.03 Parent Level Dupes + Family	Returns parent duplicates with family added. Duplicate not determined on family.
05.01 Single Recipient Emails + Family	Returns emails with one recipient and their attachments.
05.02 Bulk Emails (50+) + Family	Returns emails with 50 or more recipients and attachments.
06.01 Search Term Report	Search across all emails with text for STR.
06.01 Responsive and Family	Returns documents with Designation set to Responsive and family included.
06.02 Conflicts with Family	Returns items not coded the same as parent items.
06.03 Produce with Images Placeholders	Returns items marked with choice in Productions field, Nothing in Bates field, and Check on type of production as Images with Placeholders.
06.04 Produce Natives Only	Returns items marked with choice in Productions field, Nothing in Bates field, and Check on type of production as Native only.
06.05 Produce Natives with Placeholders	Returns items marked with choice in Productions field, Nothing in Bates field, and Check on type of production as Native with Placeholders.
By Filetype	Returns all documents with a filetype sorted by filetype.
By Folder	Returns all documents with a folder path and sorted by folder path.

11.9 Saved search history

On the History tab, you can view the audit records for saved searches, unsaved searches, and queries performed on views. For example, when a user filters a column.

Use the following guidelines to view search history records:

- **Saved Searches**—the Name column displays the name of the search. The Object Type column displays *Search*. Click the Query link to display a pop-up window with the SQL statement for the query.

Query Text

Close

```
/* <Comments>
  <ArtifactID>1003684</ArtifactID>
  <ArtifactTypeID>10</ArtifactTypeID>
  <UserID>1181490</UserID>
  <WorkspaceID>1181771</WorkspaceID>
  <QueryType>IdList</QueryType>
  <QuerySource>View or Search</QuerySource>
</Comments> */

SET NOCOUNT ON
SELECT TOP 1000
[Document].[ArtifactID]

FROM
[Document] (NOLOCK)
WHERE
[Document].[AccessControlListID_D] IN (1)
AND
((( [Document].[ExtractedText] LIKE N'%exhibit%'))
ORDER BY
[Document].[ExtractedText] ,      [Document].[ArtifactID]

-----
-- records returned: 63
-----
```

- **Unsaved Searches or Views**—the Name column displays the name of the view in which the search was performed. The Object Type column displays *View*. Click the name link to display the details page for the view. You can also click the Query link to display a pop-up window with the SQL statement for the query.

12 Search terms reports

Search terms reports provide the ability to identify documents containing specific keywords or terms. You can enter multiple terms and generate a report listing the number of hits for each term in a document. You can also select an option to create a multiple object field for the search terms report to use in your persistent highlight sets. When you select a search terms report for use with a persistent highlight set, the report determines which terms or phrases to highlight in the documents through the Review Interface.

12.1 Permissions


You may have an occasion where you want to grant limited permissions to a group of users outside of the site administrators. For example, you want to limit users from creating new search terms reports. However, at the same time, you want users to view, edit, and add to existing terms lists.

To add search terms reports permissions to a user group:


1. Create a new user group and add it to your workspace.

Note: If you do not know how to create a new user group, see The Admin Guide. If you need help adding the group to your workspace, see The Admin Guide.

2. From the Manage Workspace Permissions modal window, click **Edit Permissions**.
In the image below, the group name is **QC Reviewers**.



3. Locate the *Search Terms Reports* and *Search Terms Result* objects.
You want to edit permissions for existing terms, so you will edit the *Search Terms Result* object. The *Search Terms Reports* object grants permissions for reports overall and not the content within the report. Click one or more icons to enable the permission level. Available levels include *Disable*, *View*, *Edit*, *Delete*, *Add*, or *Manage Security*.



4. Click **Save**.
5. Click the **Manage Workspace Permissions** link.
6. Click the **Preview** button for your group.



The workspace opens with the group permissions applied. This view is helpful if you want to see how the workspace looks and functions as a group member. A banner across the top of the page reminds you of the preview mode.

You are previewing as a member of the QC Reviewers group. [Exit Preview](#)

7. Open a search terms report.
You now see buttons above the search terms list corresponding to the permissions granted. In the image below, add and edit permissions apply.
8. Check the box next to one or more search terms.
The permission level buttons become active.

The screenshot shows a web interface for managing search terms. At the top, there's a header with the word 'Terms' and two buttons, 'Add Terms' and 'Edit', which are highlighted with red rectangular boxes. Below this is a table with three columns: 'Term', 'Text Color', and 'Background Color'. The first row of the table has a checked checkbox in the 'Term' column, followed by the text '*construction:', 'Black' in the 'Text Color' column, and 'Orange' in the 'Background Color' column.

9. Click the appropriate button to edit the search term.

Note: If the only permission level granted is Edit, you will not see any buttons above the list. In this case, you can click any search term to enable a text box. Complete your edits and click Enter to save and exit.

This screenshot shows the same interface as the previous one, but the 'Add Terms' and 'Edit' buttons are no longer visible. The 'Term' column of the table now contains a text input field with the text '*construction:' inside it. This input field is highlighted with a red rectangular box, indicating it is the active element for editing.

12.2 Create and edit search term reports

You can create, edit, and copy search term reports on the Search Term Reports tab.

12.2.1 Considerations

Use the following guidelines to ensure that your search terms report properly highlights the required terms:

- Define a saved search using conditions that return the required group of documents for the **Searchable Set**.
 - Persistent highlighting applies only to documents in the searchable set.
 - Relativity only counts hits that fall within the searchable set.
 - Related items not included in the searchable set will not be searched for search terms.
- Confirm that your dtSearch index includes all documents in the **Searchable Set** of the search terms report.
- Select the **Tag Hits** toggle to create a multiple object field for the search terms report. If this field is not created you cannot select it in the **Highlight Fields** option when creating a persistent highlights set. See documentation on Persistent highlight sets.
- When using a search terms report as a highlight source in a persistent highlight set, only the terms in documents associated with the current reports appear highlighted. If you add new search terms to the reports, you must run pending terms so that they appear highlighted in documents.
- The system automatically preserves the precise order in which the terms were entered when generating a report.

Note: Fuzzy searching is not supported in Search Terms Reports (STR). You must use the Dictionary function to perform fuzzy searching. You would then add the desired terms to the STR. For information on how to use the Dictionary, see [Running a Dictionary search](#).

- We recommend running multiple queries that return small numbers of documents.
 - We recommend writing specific search terms.
 - Avoid using one broad query.
- We recommend limiting proximity operators.
 - Large numbers of proximity operators within a single query can cause slower performance. For example, *w/ or PRE/*.
 - Nesting proximity operators can also cause performance problems. For example, *dog w/5 (cat w/2 bird)*.
 - Queries can be difficult to read and understand.
- We recommend avoiding wildcards in your search terms.
 - Wildcards can slow searches.
 - Wildcards can return more documents than intended.
 - Leading wildcards should not be used. For example **ing*.

12.2.2 Create search term report

To create a new search terms report, you must add search terms report information, add terms and highlights, and run the report.

12.2.2.1 Adding search terms report information

To add STR information:

1. Navigate to the **Search Terms Reports** tab.
2. Click **New Search Terms Report**.
3. Complete the fields on the form. See [Fields](#).
4. Click **Save**.

Search Terms Reports ▾Summary ReportsPivot ProfilesDomains

Search Terms Reports ▾✎

SaveSave and NewSave and BackCancel

Search Terms Report Information

Name*

Index*

Select

Searchable Set*

Select

Include Relational Group

Select

Tag Hits

☒

Show In Field Tree

☒

Calculate Unique Hits

☐

Remove Hidden Characters

☒

Email Notification Recipients


Notes

After saving the search terms report, the Search Terms Report Status section and Search Terms Report console appear. As the status section indicates, you must add terms to your new report. See [Adding terms and highlight colors](#).

Fields

Search terms reports contain the following fields:

- **Name**—the search terms reports name. This value cannot exceed 75 characters.
- **Index**—the dtSearch index used to create the report.

 Relativity one

Searching Guide

156

- **Searchable set**—a saved search that includes a set of documents you want to use for your search terms reports.
 - The Searchable Set Index field displays the index associated with the saved search you select.
 - The index for the saved search may differ from the search index used to create the search terms report.

Search Terms Report Information

Name* Salt and Pepper STR

Index* STR Search Index **Select** [Clear](#)

Searchable Set* All Documents **Select** [Clear](#)

Searchable Set Index dtSearch

Include Relational Group **Select**

- **Include relational group**—includes the "Documents with Hits + [Group Name]" counts for each term in the search terms results.
This value counts the documents with hits for each term as well as all documents in the same relational group as the documents with hits. Include relational group only includes hits of related items in the searchable set. Relativity does not look outside of the searchable set. Click **Select** and then choose a relational group to include.
 - **Tag Hits**—if enabled, saves the results to a multiple object field named after the search terms report with the prefix STR.
 - For example, *STR - Industry terms*. This is so someone can review the results later.
 - Tags each document containing search hits using the STR multiple object field with the search terms found in each document.
 - **Show in Field Tree**—if enabled, automatically adds the terms to the field tree on the Documents view.
 - **Calculate unique hits**—if enabled, includes a unique hits value for each term in the search terms results.
 - Unique hits reflect the total number of documents returned by a particular term and only that particular term.
 - Unique hits is the count of documents in the searchable set returned by only that particular term. If more than one term returns a particular document, that document is not counted as a unique hit.
-
- Note:** Unique hits can help you identify terms in your search terms report that may be overly inclusive.
-
- **Remove Hidden Characters**—if enabled, automatically filters out hidden or non-displayable text control characters when creating or editing terms for the Search Terms Reports.
 - These hidden control characters have been known to cause issues in searching and can be unknowingly copied from Word or Excel documents.
 - The list of control characters that are filtered are the same as the default ignore section of the dtSearch alphabet file text.

- **Email notification recipients**—specifies recipients to send an email notification to when your search terms report finishes running. Enter the email addresses of the recipients. Separate entries with a semicolon.
- **Notes**—enter notes specific to the search terms report.

When viewing a list of search terms for a search terms report, term strings are limited to a set number of characters by default. Hover over a term to view the term string in its entirety.

Terms Add Terms Edit Delete			
<input type="checkbox"/> Term	Text Color	Background Color	Status
<input type="checkbox"/> *construction:	Black	Orange	Completed
<input type="checkbox"/> Candidate residency for all global regio...	Black	Light green	New
<input type="checkbox"/> Candidate residency for all global regio...	Black	Light green	New
Candidate residency for all global regions excluding zone one two and three			

12.2.2.2 Adding terms and highlight colors

To add terms to your search terms report:

Note: You must create a persistent highlight set for highlighted terms to appear in your documents. See [Persistent highlight sets](#) for more information.

1. Click the name of your search terms report.
2. Click **Add Terms**.
3. Enter your terms in the text box so that each term appears on a separate line.
Alternatively, you can also click the **Dictionary** link to display the Dictionary Search pop-up. In the pop-up you can perform searches using fuzziness, proximity, and stemming. Click **Copy to List** to add the Dictionary search results to the New Terms text box.

Note: Each line is treated as an individual dtSearch query. For more information about dtSearch, refer to the [Searching Guide](#).

4. (Optional) Select a background color and text color using the color picker. See the preview text to verify that the resulting highlighted text is readable. By default, highlighted terms appear as black text with an orange background.
5. Click **Add** to add your new terms to the terms list.

Note: A single term has a character limit of 450.

A confirmation message displays with the count of new terms added. Duplicate terms are ignored. After adding new search terms to an existing report, you must run the terms so that they appear highlighted in documents. For more information, see [Running a search terms report](#).

12.2.2.3 Workspace templates

Additionally, you can create STRs through workspace templates using the following steps:

From your active workspace:

1. Navigate to the Object Type tab.
2. Select the **Search Terms Reports** object.

3. Confirm the *Copy Instances On Workspace Creation* field is set to **Yes**.

The screenshot shows a dialog box with two sections: 'Object Type Information' and 'Object Type Settings'.

Object Type Information

Name	Search Terms Reports
Parent Object Type	Workspace

Object Type Settings

Pivot Enabled	No
Sampling Enabled	No
Lists Enabled	No
Snapshot Audit On Delete *	<input checked="" type="checkbox"/>
Use Relativity Forms	Yes
Copy Instances On Workspace Creation	Yes
Relativity Applications	Search Terms Report

4. Create a new workspace using the active workspace as your workspace template.

The search terms reports are copied over to the new workspace.

12.2.3 Editing search term reports

After running the terms, you can go back to edit terms and colors, delete terms, or copy a report.

12.2.3.1 Editing terms

To edit a term in your search terms report:

1. Click the name of your search terms report.
2. In the Term column, click on the term you would like to change.
The field will become editable.
3. Edit the text, then press **Enter**.
4. **Run Pending Terms** in order to update the search terms report. For more information, see [Running a search terms report](#).

12.2.3.2 Editing highlight colors

To edit the background and text color of terms:

1. Click the name of your search terms report.
2. Select the checkbox next to the terms you want to change, then click **Edit**.
3. Choose the new background and text color, then click **Save**.

Note: Changes to background and text color apply to all terms being edited.

12.2.3.3 Deleting terms

To remove terms from the search terms report:

1. Click the name of your search terms report.
2. Select the checkbox next to the terms you want to edit, and then click **Edit**.
3. Click **Delete**.
4. Click **Delete**.

If you remove search terms from the reports, the terms automatically disappear from the search terms report results. You must run the report again for accurate totals in the status bar and when using View Term Report.

12.2.3.4 Copying a search terms report

You can copy an existing search terms report using the mass copy operation.

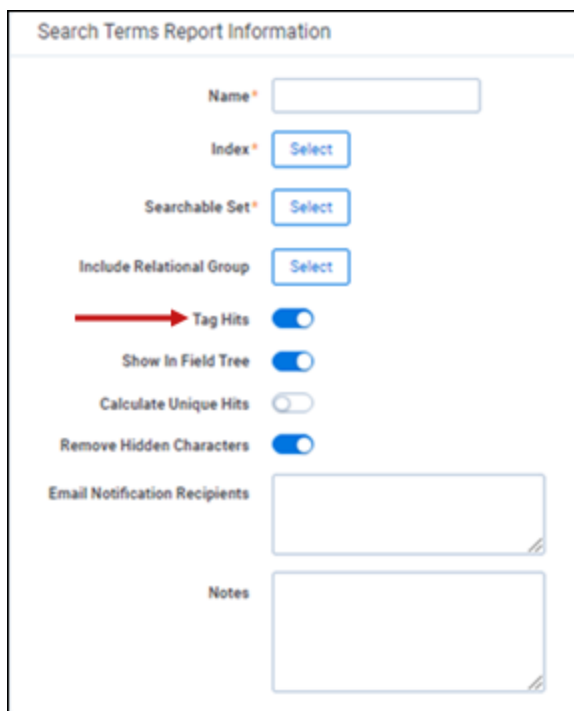
To copy a search terms report:

1. From the **Search Terms Reports** tab, click the checkbox next to the search terms report you want to copy.
2. From the mass operations bar, select **Checked**. Alternatively, you can select **All** to copy all search terms reports.
3. Select **Copy** in the drop-down menu.
The Copy window displays.
4. Click **Ok** to copy the items.

12.3 Tag Hits

When you enable the **Tag Hits** option, you can access documents in the Field Tree browser and tag terms in a saved search. Selecting the Tag Hits option creates a multiple object field for the search terms report. If this field is not created you cannot select it in the Highlight Fields option when creating a persistent highlights set. .

When enabled, Relativity saves the results to a multiple object field named after the search terms report with the prefix *STR*. For example, *STR - Industry terms*. This is so you can review the results later. During the search, Relativity tags each document containing search hits using the STR multiple object field with the search terms found in each document.

A screenshot of a web form titled "Search Terms Report Information". The form contains several input fields and toggle switches. The fields are: "Name" (text input), "Index" (dropdown menu with "Select" button), "Searchable Set" (dropdown menu with "Select" button), "Include Relational Group" (dropdown menu with "Select" button), "Email Notification Recipients" (text input), and "Notes" (text area). The toggle switches are: "Tag Hits" (indicated by a red arrow pointing to it, currently turned on), "Show In Field Tree" (currently turned on), "Calculate Unique Hits" (currently turned off), and "Remove Hidden Characters" (currently turned on).


Search Terms Report Information

Name

Index

Searchable Set

Include Relational Group

 Tag Hits ☒

Show In Field Tree ☒

Calculate Unique Hits ☐

Remove Hidden Characters ☒

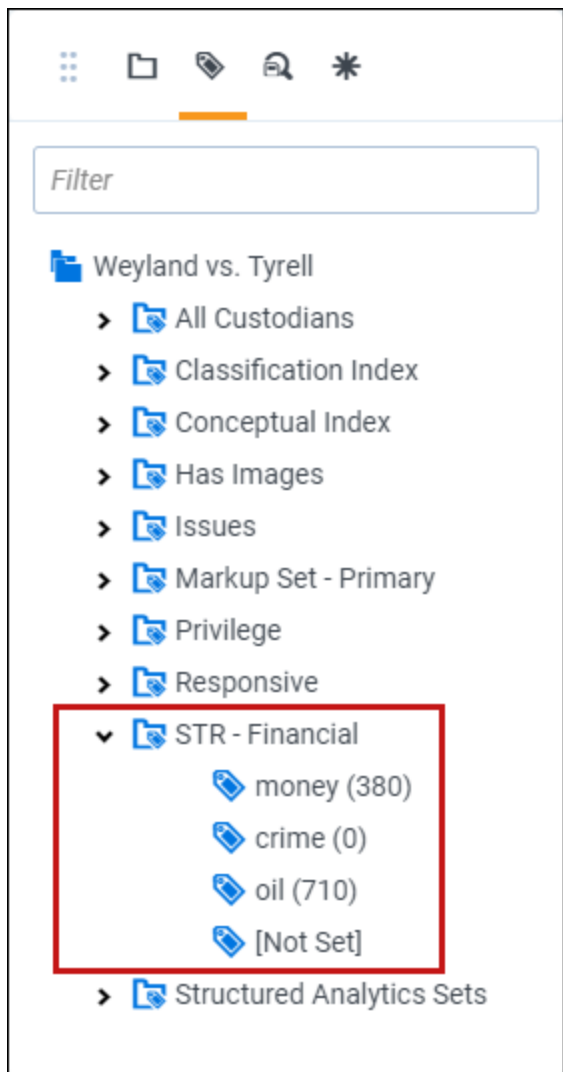
Email Notification Recipients

Notes

12.3.1 Tagged documents in the Field Tree browser

After you run a search terms report with the **Tag** option enabled, Relativity creates a folder in the Field Tree browser with documents grouped by tags found by the search terms report. Each tag includes the count of documents containing that term. The folder is named after the STR multiple object field created by the search terms report. For example, if your STR is named Produced Documents, the folder in the Field Tree browser is named STR - Produced Documents.

Click a search term tag in the Field Tree browser to view documents in your searchable set tagged with the selected term. You can also email a link to the tagged documents by right-clicking the tag results and selecting **Email Link**.



12.3.2 Tagged search terms in a saved search

After you run a search terms report with the **Tag Hits** option enabled, Relativity creates choices for each of the terms that you specified. You can then use these choices as criteria in a saved search.

Use this procedure to create a saved search using tagged search terms:

1. Follow the instructions for setting fields in the Information and Search Conditions sections on a saved search.
2. Click on the Fields tab in the Saved Search pop-up.
3. Select the desired field in the Unselected section and move it to the Selected section using the arrow icons.
4. Click the **Condition** button and select an option from the pop-up.
5. Perform the following tasks on the **Select STR** dialog:
 - a. Select your search terms report in the Field box.
 - b. Select an Operator, such as *any of these*.
 - c. Select the STR option from the **Condition** drop-down menu to select search terms on the **Select Items** dialog.
 - d. Set any other fields as necessary.

6. Click **OK**.

For example, if you added the tagged search terms called money, crime, and oil, the Conditions section would appear as follows:

The screenshot shows a dialog box titled "Condition: STR - Financial". At the top, there is a dropdown menu set to "these conditions". Below it is a "Condition" dropdown and a "Logic Group" button. The main area contains three conditions, each in a white box with a light blue border, connected by "AND" logic. The first condition is "1. STR - Financial" with "any of these: money". The second condition is "2. STR - Financial" with "any of these: crime". The third condition is "3. STR - Financial" with "any of these: oil". At the bottom right, there are "Apply" and "Cancel" buttons.

7. Click **Apply** and **Save & Search** to run your query.

Repeat steps 1-3 for each tagged search terms.

12.4 Run a search terms report

You generate a search terms report by using the options available in the search terms report console. The console appears after you save a search terms report or when you open an existing report from the Search Terms Reports tab.

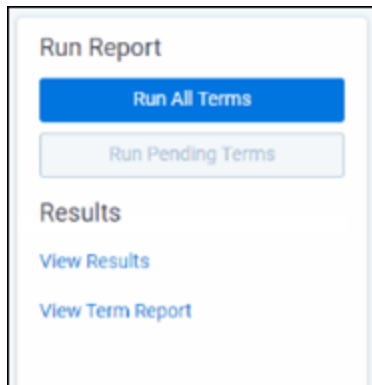
The console includes the following options:

- **Run All Terms**—generates counts for each term. Use this option when generating the report for the first time or if you want to regenerate counts for all terms in the report. Run all terms after adding new documents to the searchable set.
- **Run Pending Terms**—updates an existing report. It runs a report on only those terms with a Pending status.
- **View Results**—opens the *Search Terms Results (All)* page. This page displays the report results, listing the number of document hits for each term. From the Search Terms Results (All) page, click the *To Report* button

to return to this page.

 To Report (STRNumbers)

- **View Term Report**—opens the report in a graphical format. You can print or save the report. To save the report, click the *Export To* drop-down menu at the top of the page and select a file type, such as PDF. The file export begins automatically and saves the file to your local download location.



The Search Terms Reports > Terms table has the following columns:

- **Name**—search term included in search terms report.
- **Status**—shows the current status of the searched term. Options include Completed, Pending, and Error.
- **Documents with hits**—the number of documents in the searchable set that contain the search term.

Note: Documents with hits is not security-aware or influenced by permissions. This means that it includes documents the user cannot view in a basic search. For example, a user could perform a dtSearch that returns a total of five documents, including two inaccessible documents. Even though the user can only view three documents, the search terms count still includes all five documents originally tagged with the search term.

- **Documents with hits, including group**—counts the documents with hits for each term as well as all documents in the same relational group as the documents with hits. The count only includes hits of related items in the searchable set. It will not look outside of the searchable set.
- **Unique hits**—counts the number of documents in the searchable set returned by only that particular term. If more than one term returns a particular document, that document is not counted as a unique hit. Unique hits reflect the total number of documents returned by a particular term and only that particular term.
- **Last Run Time**—time stamp when the search terms report last ran.
- **Last Updated**—time stamp when the search term was last updated.
- **Error messages**—displays error messages for report terms, if encountered. See [Managing errors](#) for more details on error handling in search terms reports.
- **Text Color**—the font color you set when adding the term.
- **Background Color**—the background color you set when adding the term.

In this page, you can also access a list of any terms that failed during the creation of the search terms report. To read these error messages, change your view to *Search Terms Results (All)*.

Note: RelativityOne now runs up to 20 search terms reports in parallel, simultaneously, on a single instance across any number of workspaces. You can, however, queue as many reports as needed and they will run when others complete.

12.4.0.1 Search terms reports status

After running a search terms report, the search terms reports status section appears. It lists the search terms report name and status. The status indicates the current progress of the report. This field contains either Searching your terms, Completed, or Error.

Report Status: Completed			
Number of Terms	Documents with Hits	Documents with Hits + Relational Groups	Documents in Searchable Set
2	380	0	15937

This section also provides a summary of the search terms report and its results:

- **Number of Terms**—total number of terms run in the search terms report.
- **Documents with Hits**—the number of documents in the searchable set that contain the search term.
- **Documents with Hits + Relational Groups**—counts the documents with hits for each term as well as all documents in the same relational group as the documents with hits.
- **Documents in Searchable Set**—total documents in the designated searchable set.

12.5 Search terms report hit count

This topic describes how to calculate and store the number of terms in a Search Terms Report (STR) that hit on a document.

You can then sort, filter and search for documents based on the number of STR hits, allowing you to prioritize review based on STR hit count or create searches that return documents having a minimum number of STR hits

Note: This workflow requires downloading and installing the **Delimiter Count by Saved Search** script. See [NSerio / delimiter-count-by-saved-search](#) on GitHub for a link to the script.

12.5.1 Prerequisites

Before calculating and storing hits, you must complete the following steps:

1. Download the **Delimiter Count by Saved Search** script. See [NSerio / delimiter-count-by-saved-search](#) on GitHub for a link to the script.
2. Create two STR-specific fields.

12.5.1.1 Delimiter Count by Saved Search script

Your workspace should have this script in the script library.

Note: If not, download and install the **Delimiter Count by Saved Search** script. See [NSerio / delimiter-count-by-saved-search](#) on GitHub for a link to the script. For more information on adding a script from the library, see Scripts documentation.

Add the script to your workspace:

1. Navigate to and select the **Scripts** tab.
2. Click **New Relativity Script**.
3. Choose **Select from Script Library** on the **Script Type** field.
4. In the **Relativity Script** field, click **Select**, and choose the library script you want to add.
5. Click **Save**.

12.5.1.2 Search term report fields

You must create two STR-specific fields:

- STR Hit Terms
- STR Hit Count

When creating the **STR Hit Terms**, add the following information to the fields:

- **Name**—STR Hit Terms
- **Object Type**—Document
- **Field Type**—Long-text field

When creating the **STR Hit Count**, add the following information to the fields:

- **Name**—STR Hit Count
- **Object Type**—Document
- **Field Type**—Whole Number

12.5.2 Saved search

The next step in the process is to create a saved search that returns the documents that hit on one or more terms for a specific STR.

The STR used must have the Type field set to **Report and Tag**. For details on creating a search terms report, see the Admin Guide. Ensure that the saved search includes the Control Number (Identifier) field and STR field.

Then export the results of the saved search. To export the saved search:

1. Perform an **Export to File** mass operation on the results.
2. Select **Comma Separated Values (.csv)** for the export *Format*.

The STR terms will be populated in the file as semicolon delimited values.

12.5.3 Import

Now you must use Import/Export to export the saved search. For more information, see Import RDO data in the Import/Export guide.

To import the saved search:

1. Navigate to the **Import/Export Job** page.
2. Click the **New Import/Export Job** button.
3. On the *Choose Workflow* dialog, select the **Import** button.
Import should be selected by default.
4. Click the **Structured (Processed) Data** workflow box.

5. On the *Choose the Load File Type* dialog, click the **RDO** box.
6. On the *Choose Load File and Location* dialog, complete the fields, and then click **Continue**.
7. On the *Load File Settings* dialog, select the delimiters and encoding values of the load file, and then click **Continue**.

The File Column Headers Preview pane will help you to choose the proper settings.

8. On the Fields Mapping dialog, select the load file fields to import and map them to the workspace fields, then click **Continue**. Click **Auto Map Fields** to automatically map fields. Relativity then maps the fields from the load file to existing fields in the workspace of the same name.

Note: Any fields without a match will display *Choose...* in the Workspace Fields column. You must select one from the drop-down list.

9. On the Import Settings dialog, configure your import job by selecting from which line number you want the load file to begin importing and how you want the records to be imported.
10. Click the **Import** button to start the importing process.
11. Click **Close** when the importing process is complete.

Caution: If you used Import/Export in web only mode, meaning Express Transfer was not active, do not close the Importing dialog until after the Uploading step has successfully completed. Closing the dialog prior to the Uploading step completing will terminate the import process.

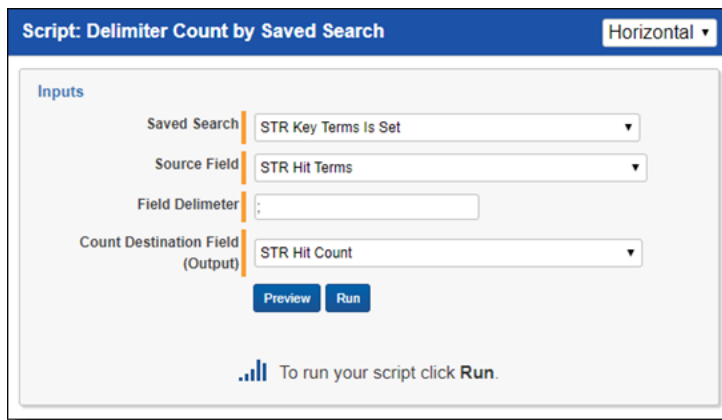
12. Click **Finish** to close the dialog.

12.5.4 Run script

Finally, you must run the Delimiter Count by Saved Search script.

1. Navigate to the **Scripts** tab in the workspace.
2. Locate and select the **Delimiter Count by Saved Search** script.
3. ClickRun the script.
4. Set the fields in the script pop-up:
 - a. **Saved Search** – The search created above that returns documents where the STR is set.
 - b. **Source Field** – The **STR Hit Terms** field that was updated above via Import/Export import.
 - c. **Field Delimiter** – Enter a semicolon (;).

- d. **Count Destination Field** – The **STR Hit Count** field created above. Note: this field is only populated when its value is null/blank; it will not be overwritten by the script.

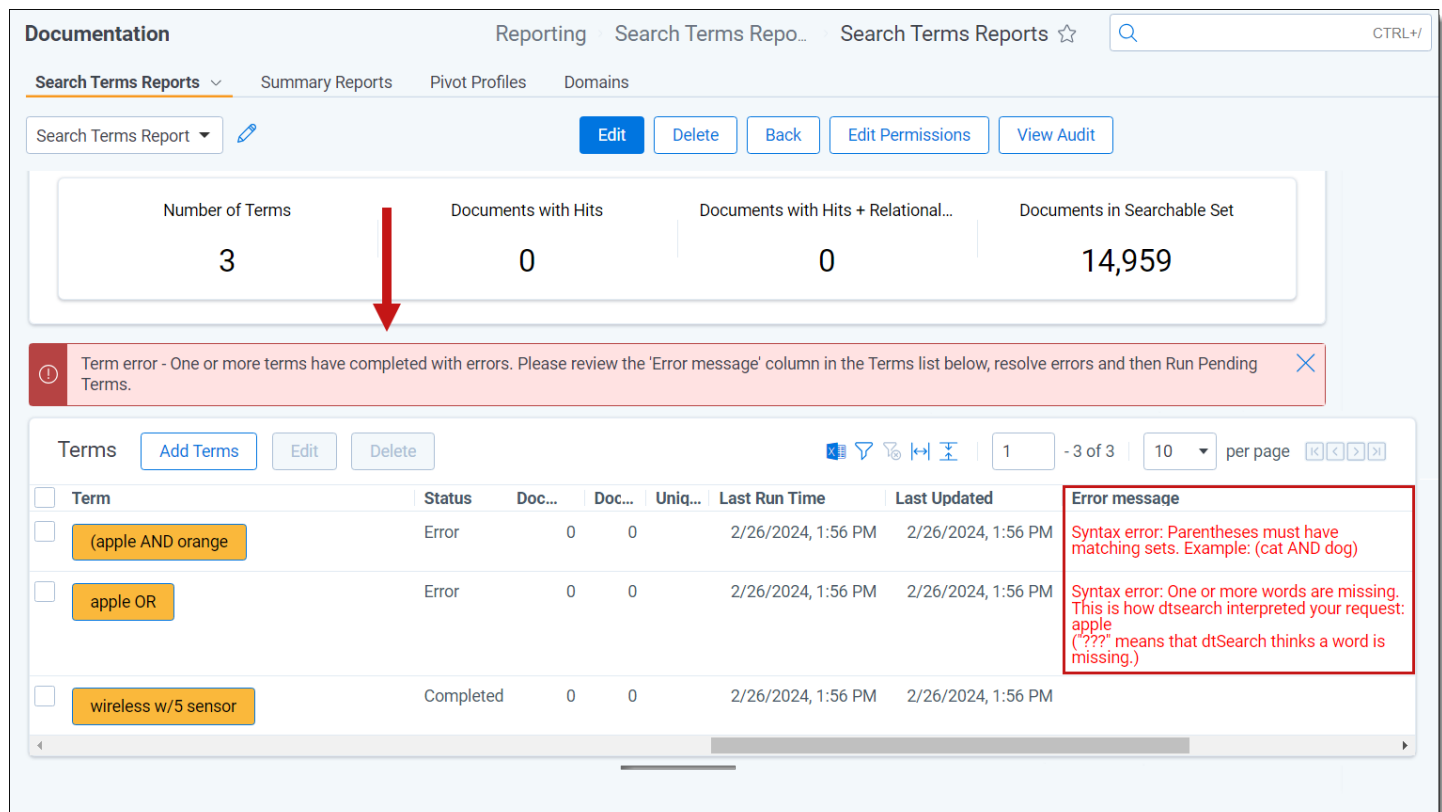


5. In the *Actions* console, click **Run**.

The **STR Hit Count** field now contains the number of STR terms hitting on each document. You can use the STR Hit Count field in searches and views to sort and filter documents based on the STR Hit Count value.

12.6 Errors

As the search terms report runs, Relativity checks for errors. If Relativity encounters errors, you will see a red banner just above the list of terms. You will also see a message in the **Error message** column for each term that Relativity could not process. The **Status** column displays Error.



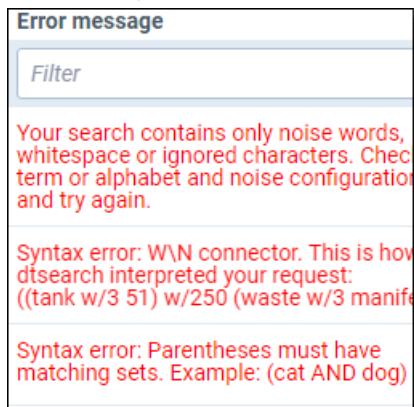
Term	Status	Doc...	Doc...	Uniq...	Last Run Time	Last Updated	Error message
(apple AND orange)	Error	0	0		2/26/2024, 1:56 PM	2/26/2024, 1:56 PM	Syntax error: Parentheses must have matching sets. Example: (cat AND dog)
apple OR	Error	0	0		2/26/2024, 1:56 PM	2/26/2024, 1:56 PM	Syntax error: One or more words are missing. This is how dtsearch interpreted your request: apple ("???" means that dtSearch thinks a word is missing.)
wireless w/5 sensor	Completed	0	0		2/26/2024, 1:56 PM	2/26/2024, 1:56 PM	

The error messages tell you the type of error encountered. For example, you can resolve syntax errors yourself and then retry them, while support errors require customer support for resolution.

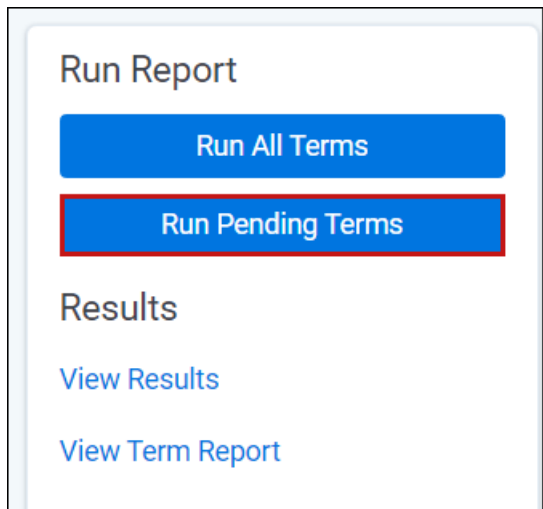
12.6.1 Syntax errors

For syntax errors, the search index could not search for the term due to an error in the syntax. The message tells you what is wrong with the syntax and how to fix it.

You can fix these errors yourself by editing the term and addressing the issue outlined in the error message. For example, the second syntax error in the image below indicates a missing parenthesis. You can edit the term and add the missing parenthesis in this case.



After you address syntax errors, you must rerun the report to capture the terms. Click the **Run Pending Terms** button in the **Run Report** console to run the report on the pending terms. When complete, the page updates with the new term status.



12.6.2 Support errors

Support errors are errors that you cannot resolve yourself. Support errors require intervention from customer support. The error message includes an expandable section with explicit details about the error, which support can use to help determine and resolve the issue.

An error has occurred. Please contact support. [▼ View Error](#)
The dtSearch Engine could not access the provided index: \$E 0001 Unable to access index
\\dtafstus014128t0160gl8s.file.core.windows.net\dtsearch001\dt_1597752_321f062c-c09e-4d63-b8d5-b414ecbc9182\Index_b3edf575-9e77-4a28-817c-4aaa9929277b\ Operation failed: open (could not open directory)
path=\\dtafstus014128t0160gl8s.file.core.windows.net\dtsearch001\dt_1597752_321f062c-c09e-4d63-b8d5-b414ecbc9182\Index_b3edf575-9e77-4a28-817c-4aaa9929277b\ (error = The system cannot find the file specified. (2)).

13 Short message search (Advanced Access)

Short message search allows for a detailed search on RSMF documents by examining message-level metadata. It helps to, for example, identify conversations between specific participants, detect messages that were deleted or edited, track message reactions, and locate messages sent within a specified time period. Once messages are found, message-level coding can be performed on them in the Viewer if necessary.

This workflow uses Elasticsearch to create the short message search index within the workspace. The index must be created before conducting searches. See [Short message search index](#) for more information.

13.1 Considerations

Before you run a short message search, consider the following regarding message-level conditions:

- Are treated as an index search condition.
- Cannot be combined with other index search conditions.
- Can only be searched with other field conditions using the AND operator.
- Can only be used as the first condition within a search. From here, you can run your search, and it will return RSMF documents matching the supplied search criteria.

13.2 Creating a short message search

After you build the short message search index, you can search it from the Documents tab. See [Short message search index](#) for more information on building the index.

To create a short message search:

1. Navigate to the **Documents** tab.
2. Click **Add Condition** in the search panel to add a condition. For more information on the search panel, see [Search panel on page 32](#).
3. Select **(Short Message Search)** from the Add Condition fields list.
The Add Condition - Select Field modal appears.
4. Select the desired short message field from the list. See the [Searchable fields on page 178](#) section for a list of all fields available for short message searches.
The Filter modal appears. The data entry fields on this modal vary depending on the short message field previously selected.
5. On the Filter modal, choose the operator, enter or select query information, and click **Apply**.
6. Optionally, you can:
 - Add additional conditions by clicking **Add Condition** and repeating the steps above.
 - Create logic groups to group conditions together by clicking the Add Condition drop-down arrow and selecting **Logic Group**. Then, drag the conditions into the logic group frames. Use the AND or OR operators to join logic groups.

Short Message Search Conditions

+ Condition

▼

(1 AND 2) AND 3

1. Message Body

is dogs AND cats

AND ▼

2. Deleted

is No

AND ▼

3. Is Edit

is Yes



Clear All Conditions

Apply

Cancel

- When finished entering search conditions, click **Apply** on the Short Message Search Conditions modal. The Short Message Search condition card is added to the Search Conditions list.
- Repeat the steps above to add more conditions or logic groups.
- Click **Run Search** to start the search if the Auto Run toggle is not enabled. If the toggle is enabled, the search begins automatically.
After running a search, the Document list displays all the documents with at least one message or event matching the query conditions.
- Click the document **Control Number** in the Documents list to view it in the Viewer.

Messages containing a highlight are identified in the Viewer with a yellow highlight around the metadata  icon.

Navigate to each one by using the Persistent Highlight navigation buttons ( ) on the top toolbar of the Viewer. For more information on the Viewer, see Short Message Viewer and Short Message coding.

If needed, you can create a saved search on the Documents list as you would any other search.

13.3 Defining search criteria operators

There are five field types in short message search conditions:

- [Long text operators below](#)
- [Fixed-length text operators on page 175](#)
- [Date operators on page 176](#)
- [Whole number operators on page 177](#)
- [Yes/No operators on page 177](#)


Each field type works similarly to the standard Relativity field type of the same name, as defined below. The primary exception is long text, which uses advanced text search syntax that is similar to dtSearch. See [Long text operators below](#) for more information.

See the [Searchable fields on page 178](#) section for a list of all fields available for short message searches.

13.3.1 Long text operators

Use text queries to search for long text. This enables searching Elasticsearch indexes with a syntax similar to dtSearch.

Note: Elasticsearch index does not support all dtSearch syntax.



The [advanced text search syntax](#) is used in Long Text condition values.

Note that all phrases must be contained within quotation marks. Please see the linked documentation for further details.

Filter: Message Body

Operator

advanced text search ▼

Enter text query

Apply

Cancel

The table below shows search string examples and their expected results. You may also search using the “Is set” or “Is not set” operators to check for empty or filled values.

Search string	Returns RSMF documents with messages containing...
"apple pear"	The exact phrase <i>apple pear</i> .
apple AND pear	The words <i>apple</i> and <i>pear</i> .
apple OR pear	Either <i>apple</i> or <i>pear</i> .
apple W/5 pear	The word <i>apple</i> within five words of <i>pear</i> . The opposite is also true.
apple PRE/5 pear	The word <i>apple</i> within five words before <i>pear</i> . The opposite is not true.
apple AND NOT (pear)	The word <i>apple</i> but not the word <i>pear</i> .
pear*	Any word beginning with <i>pear</i> , including the word itself. For example, <i>pears</i> and <i>Pearson</i> .
pe?r	Any word beginning with <i>pe</i> with a single character between that and <i>r</i> . For example, <i>pear</i> or <i>peer</i> .

13.3.1.1 Reserved characters and words

Similar to dtSearch, certain characters and words are reserved by the advanced text search syntax and require special handling to be searched.

- **Reserved characters:** " () % * ? ~ ! / \ + - & | { } [] ^ : < >
- **Reserved words:** AND, OR, NOT, W/N, PRE/N (where "N" is a number)

To search for:

- **Reserved characters**—insert a backslash (\) before them. Examples:
 - To search for *The cat is in the hat!*, enter `The cat is in the hat\!`.
 - To search for the literal phrase, *The query should be W/5*, enter `The query should be W\ /5`. That is **W** followed by backslash (\), then forward slash (/), and finally **5**.
- **Reserved words**—use quotation marks around them, either individually or as part of a phrase. Example query: *"Bob and Jill went to the market."*
 - The word *and* is automatically searched because, logically, the AND operator cannot be used within a phrase.
 - To search for the word *and* separately, enter the query as: `bob OR "and" AND jill`

13.3.1.2 New syntax rules

As mentioned, the advanced text search syntax and dtSearch differ in their query requirements. Certain queries need specific criteria to be valid with the advanced syntax. These are detailed below.

- Enclose all phrases within quotation marks. If not, a syntax error results.
- Any query containing multiple Boolean clauses must have each clause contained within parenthesis to ensure that the intended meaning of the query is clear.
 - Valid: `(apple AND pear) OR (dog AND cat)`
 - Invalid: `apple AND pear OR dog AND cat`
- Follow the NOT operator with parenthesis containing the content to be searched. This clearly identifies the content the NOT operator is acting on.
 - Valid: `apple AND NOT (pear)`
 - Invalid: `apple AND NOT pear`

13.3.1.3 Proximity operator (W/N, PRE/N) limitations

Compared with dtSearch, there are a few limitations described below on how you can use the two proximity operators.

- Do not use AND with proximity operators.
 - Invalid: (cat AND dog) W/7 bird
- Do not nest proximity operators.
 - Invalid: (cat W/3 dog) W/5 bird
- NOT cannot be used with W/N or PRE/N.
 - Invalid: cat NOT W/5 dog

13.3.1.4 Analyzer and wildcard behavior

Relativity's short message search indexing uses Elasticsearch's standard analyzer. This analyzer determines the rules for segmenting words into tokens, similar to the function of dtSearch Alphabet and Noise Word files. Unlike those configurations, it does not require any customization within RelativityOne.

The standard analyzer follows Unicode rules for word breaking and punctuation, typically removing punctuation marks. For instance, "bite-sized" would be processed as "bite sized," with the word being treated as two separate tokens. There are edge cases where punctuation is indexed, especially with uncommon punctuation or foreign language text. However, the underlying rules are not consistent enough to be reliably summarized or predicted.

This may lead to overly inclusive results in the case of words with punctuation in them.

- For example, "bite\ sized" (escaping the hyphen character with a backslash) will match both "bite-sized" and "bite sized." The hyphens are treated as whitespace in both the search term and the indexed text.

The behavior differs when using wildcards (* or ?), as search terms with wildcards are not analyzed. Search terms containing punctuation, such as hyphens, are processed exactly as they appear, with punctuation included in the search.

- For example, "bite\ size*" (with the hyphen character escaped using a backslash) will not match any entries because the escaped hyphen is searched for but is not in the index.
- There are multiple options for working around this behavior:
 - One is to search for the word both with and without punctuation. This will cover normal cases, plus any potential edge cases where the punctuation is present in the index. Example: bite\ size* AND "bite size"
 - Another is to use the proximity operator. Example: bite W/2 size*

The table below provides examples:

Search Term	Analyzed Term	Text in Document	Analyzed Text in Index	Match?
bite\ sized	bite sized	bite-sized	bite sized	Yes
bite sized	bite sized	bite-sized	bite sized	Yes
bite\ size*	bite\ size*	bite-sized	bite sized	No
bite W/2 size*	bite W/2 size*	bite-sized	bite sized	Yes

13.3.2 Fixed-length text operators

Fixed-length text values are entered using a simplified syntax relative to that of long text operators. Both single character wildcards (?) and multiple character wildcards (*) are supported, as well as escaping those characters using the backslash (\).

- **Single character example:** *pa?s* would match on both *pads* and *pass* when using the “Matches” operator.
- **Multiple character example:** *d*ing* would match on *daring*, *dashing*, and *dancing*.
- **Escaping example:** *Did the query work\?* would match that sentence including the question mark.

Fixed-length fields support the following operators: Matches, Does not match, Is set, and Is not set.

i

- Fixed-Length condition values support the use of single character wildcards (?) and multiple character wildcards (*).
- The backslash character (\) can be used to escape ? and * characters in order to search for them.

Filter: Reaction Types

Operator

matches ▼

Enter text query

Apply

Cancel

13.3.3 Date operators

The following operators are available for date field types. Values should be entered in UTC rather than your local time zone.

Filter: Timestamp

is ▼

mm/dd/yyyy

hh:mm tt

mm/dd/yyyy

hh:mm tt

Apply

Cancel

Operator	Returns Documents Where...
Is	The entered date is equal to the field value.
Is not	The entered date is not equal to the field value.
Is set	The field is not empty.
Is not set	The field is empty.
Is before	The field value is before the entered date.
Is before or on	The field value is before or on the entered date.
Is after	The field value is after the entered date.
Is after or on	The field value is on or after the entered date.
Between	The field value is between the two entered dates.

13.3.4 Whole number operators

The following operators are available for whole number field types.

Filter: Artifact ID

Operator
Select Operator ▼

Enter a number

Enter a number

+ Add condition

Apply

Cancel

Operator	Returns Documents Where...
Is	The entered number is equal to the field value.
Is not	The entered number is not equal to the field value.
Is set	The field is not empty.
Is not set	The field is empty.
Is less than	The field value is less than the entered number.
Is greater than	The field value is greater than the entered number.

13.3.5 Yes/No operators

The following operators are available for Yes/No field types.

Filter: Responsive

Is ▼

☐ Yes

☒ No

+ Add condition

Apply

Cancel

Operator	Returns Documents Where...
Is	The selected value (Yes or No) is equal to the field value.
Is not	The selected value (Yes or No) is not equal to the field value. This operator covers situations where the field is either the opposite of the chosen value or not set. For example, <i>Is not</i> with a value of Yes would cover values of No or unset values.
Is set	The field is not empty.
Is not set	The field is empty.

13.4 Searchable fields

All searchable fields for short message searching.

Field name	Type	Explanation	Example
Attachment Count	Integer	Number of attachments on a message.	2
Attachment	Fixed-Length Text	Attachment filenames, including file extensions.	Sample.txt
Attachment IDs	Fixed-Length Text	Internal IDs referring to the attachments – generally matches the filename.	Sample.txt
Attachment Max Size	Integer	Max size of any attachments on a message.	381734829
Conversation Display	Long Text	Overview field that displays either the message platform or the channel and users.	Platform: Microsoft Teams Channel Name: new-project
Conversation Platform	Fixed-Length Text	Platform where the conversation took place.	Slack

Field name	Type	Explanation	Example
Conversation ID	Fixed-Length Text	Internal ID referring to the conversation.	2386477
Conversation Type	Fixed-Length Text	Type of conversation, generally direct or channel.	direct
Deleted	Yes/No	Whether or not the message was deleted.	Yes
Document ID	Fixed-Length Text	Internal ID of a document.	1387384
Edited	Yes/No	If the message was edited at any time, it returns Yes for any version of the edited message, including the original.	Yes
Event ID	Fixed-Length Text	Unique internal identifier for the event in the RSMF manifest.	C1
Event Important	Fixed-Length Text	Message importance, generally normal or high.	Normal
Event Parent	Fixed-Length Text	Unique internal identifier of the parent of the RSMF event, such as the thread that a message belongs to.	B1
Event Type	Fixed-Length Text	Type of event, generally message, join, disclaimer, or leave.	message
Message Body	Long Text	The body of the individual message.	This is an example of message body text.
Notes	Long Text	Message-level coding notes made in the Notes field of the Coding Layout.	This is an example of note text.
Privileged	Yes/No	Whether or not a message was coded as Privileged in the message-level coding layout.	Yes
Reaction Participants Display	Long Text	Names of conversation participants who left reactions on the message.	John Smith
Reaction Total Count Across Types	Integer	Total reactions left across a message.	5
Reaction Types	Fixed-Length Text	Type of reactions left on a messages.	slightly_smiling_face
Reaction Types Count	Integer	Number of reaction types left on a message.	3
Responsive	Yes/No	Whether or not a message was coded as Responsive in the message-level coding layout.	Yes

Field name	Type	Explanation	Example
Sender Account ID	Fixed-Length Text	Primary ID distinguishing the sender in the application, such as a phone number or social media handle.	1-555-921-4560
Sender Display	Long Text	Display name of the sender.	John Smith
Sender Email	Fixed-Length Text	Email address of the sender, if applicable to that platform.	john.smith@gmail.com
Sender ID	Fixed-Length Text	Unique internal identifier for the sender in the RSMF manifest.	P1
Timestamp	Date	Date and time that the message was sent.	5/1/2025 1:00 AM

13.5 Short message search index (Advanced Access)

Short message search index enables you to perform more complex searches on RSMF documents processed in Relativity. It provides enhanced accuracy by storing and searching on the message-level and event-level metadata by using Elasticsearch. For more information on searching short message metadata after building a search index, see [Short message search \(Advanced Access\) on page 171](#).


While Elasticsearch is similar to dtSearch and filtering, there are differences:

	Description	Example	Search precision
Elasticsearch	Searches on message-level and event-level metadata.	If you enter a keyword without writing a field-specific query, Relativity searches all message-level fields, including the message body. Searching for "John" returns results that match that term everywhere Relativity finds it. Such as in the Sender Display Name and Message Body fields.	High
dtSearch	Searches cannot search specific aspects.	If you searched "John Smith," the results would show that the phrase "John Smith" appears in the document, but not if John specifically sent a message, reacted, left the chat, and so forth.	Medium
Filter	Searches on document-level fields only.	If you want to find documents that have a message sent by John Smith, the closest you can get is by checking the Participants field. The Participants field tells you if John Smith was in the DM or channel at any point in time.	Low

13.5.1 Installing the application

To start searching on short message metadata, you must first install the Search AI app to your workspace.

1. Navigate to the workspace where you want to install the application.
2. Navigate to the **Relativity Applications** tab.
3. Click **New Relativity Application** to display an application form.

4. Click the **Select from Application Library** radio button in the Application Type section.
5. Click the ellipses  icon next to the **Choose from Application Library** field.
6. Select **Search AI** in the Select Library Application list. This list only displays applications added to the Application Library. If Search AI is not included in the list, see the Installing applications topic.
7. Click **Ok** to display the application in the **Choose from Application Library** field. The application form also displays the following fields:
 - **Schema Version**—displays the version of the application that you are installing.
 - **User-friendly URL**—displays a user-friendly version of the application's URL. This field may be blank.
 - **Application Artifacts**—displays object types and other application components by clicking the plus icon.
 - **Map Fields**—there are no fields available in Search AI for mapping.
8. Click **Import** to install the Search AI app into the workspace.
9. Review the import status of the application. Verify that the installation was successful or resolve errors.

The Search AI app is now in your workspace. You can begin building an index.

13.5.2 Short message search index workflow

The short message search index workflow is similar to other search indexes. You must first build an index before being able to search it. For more information on searching short messages, see [Short message search \(Advanced Access\) on page 171](#).

Although Elasticsearch works with messages and events, it returns results at the document level.

Note: Documents returned may contain one or more hits at the message or event level.

13.5.2.1 Building the full index

To build the full index, perform the following steps:

Note: All RSMF documents processed in Relativity before to May 8, 2025 must be republished before building the full index.

1. Navigate to the **Search Indexes** tab.
2. Select **All Search Indexes** from the View drop-down next to the New Search Index button.
3. Select the **Elasticsearch** entry.
4. Click **Build Full Index** located in the Build and Manage Index console.

The Status field shows *Completed* when the build finishes successfully. The build process includes all RSMF documents in the workspace, so you do not have to select a saved search or specific indexed fields.

Below are the possible statuses when building a short message index:

- **In progress**—states that the index is currently building.
- **Completed**—the build finished successfully.
- **Failed**—indicates 50% or more of the documents failed. Review the Retry Errors report for details. See [Retry errors on page 183](#) for more information.

- **Completed, Eventing active with errors**—indicates an error on a message-level coding decision. Resolve the error before continuing. See [Retry errors on the next page](#) for more information.

After the full index build is complete, you can run incremental index builds as needed instead of performing full builds. For more information, see [Running incremental index builds below](#).

Refer to [Creating a short message search on page 171](#) to start searching on the index results from the Documents list.

13.5.2.2 Running incremental index builds


After running a successful full index build, you can run incremental builds. Incremental builds only run on new RSMF documents processed into the workspace and will add them to the index. As a result, incremental builds run faster than a full build.

Consider the following items when running an incremental build:

- Incremental builds are only possible when the index status is *Active*.
- Incremental builds do not index any changes made to existing document data since the last full build.
- Canceling an incremental build before it finishes deletes the existing index data. As a result, you must run a full index before rebuilding.

To start an incremental build:

1. Click **Build Incremental Index** located in the Build and Manage Index console.
2. In the Incremental build confirmation window, select the checkbox if you want to activate the index once it's complete. Otherwise, leave it blank.



3. Click **Yes** to begin the incremental build.

The Index Manager Agent will start indexing the job as soon as it becomes available.

Upon successful completion, the newly indexed files merge with the existing search index. If the job is unsuccessful, click **Retry Errors** located in the Errors and Details console. See [Retry errors on the next page](#).

13.5.2.3 Automating incremental index builds using Automated Workflows

Automated workflows can trigger an incremental index build when the system adds new documents into a workspace.

To set up the Automated Workflow:

1. Navigate to the **Automated Workflows** tab.
2. Click **New Automated Workflow**.
3. Select **New Documents Added** as the trigger.
4. Click **Add Action** and select **Build Elasticsearch Index**.
5. Select the **Elasticsearch** index from the drop-down menu.
6. Add a **Name** and **Email** for notification on the left panel.

After creating the automated workflow, the building of Elasticsearch indexes is now automated.

13.5.3 Retry errors

To view and retry index build errors, use the options in the Errors and Details console:

- **Retry Errors**—click to retry any documents that failed to index during the earlier build.
- **Show Document Errors**—click to see individual files and their errors on the Errored Documents modal. Click the **Document Artifact ID** link to display the document in the Viewer. Click the **Export** link to download a .csv file with the error details so that you can fix them and rebuild the index. See [Retry errors workflow below](#) for more information.

Errored Documents			Export
Control Number	Document Artifact IDs	Error Category	
REL_R10000008851	2392572	Coding Error - please retry errors	
REL_R10000008852	2392576	Empty File - please re-process document	

- **View Documentation**—click to navigate to the RelativityOne documentation site.

Retry errors workflow

When errors occur, you can first try to rebuild the index with the errored documents using the Retry Errors option in the Errors and Details console. This may resolve any system issues that may have occurred, allowing the errored documents to index successfully.

If there are still errors, click the **Show document errors** link to review each document and error. Error categories include:

- **Indexing Error**—a problem occurred while indexing documents. Click **Retry Errors** to re-index the failed documents.
- **Coding Error**—a problem occurred when updating a message-level coding field. Click **Retry Errors** to re-index the failed documents.
- **Metadata Storage Errors**—a problem occurred when trying to index the metadata. Click **Retry Errors** to re-index the failed documents.
- **Empty File**—the file has no data in it. The document file did not process correctly during Processing or an error occurred during data ingestion through Integration Points. Reprocess the document in Processing or Integration Points.

You can export the document information to a .csv file using the Export link. Then, either create a saved search to exclude the documents or proceed with reprocessing them.

13.5.3.1 Eventing

When you change the fields below using message-level coding in the Viewer, the system triggers an event. Their data is automatically refreshed in the index.

- **Responsive**—this field reflects any message-level coding decisions made on the *Responsive* field using the coding layout.
- **Privileged**—this field reflects any message-level coding decisions made on the *Privileged* field using the coding layout.
- **Notes**—this field reflects any message-level coding notes made in the *Notes* field using the coding layout.

Events are searchable shortly after the "Last updated on" time, so there is no need to build an index.

If an event fails to index, the index status becomes *Completed, Eventing active with errors*. Click the **Show Document Errors** link to review the documents in the Errored Documents modal. Then, click **Retry Errors** to start fixing the issue. For more information, see [Retry errors on the previous page](#).

Also, if one or more message-level coding decisions did not index, a banner displays on the Documents list page. Navigate to the index details page to retry the documents.

14 Search Q&A (Advanced Access)

Search Q&A is a new application where you can use natural language and conversational capabilities to investigate and better understand your data, reducing the time required to find relevant information. This is an experimental capability that we are providing as part of this Advanced Access (AA) program so that Relativity can incorporate your feedback into this feature's development.

Advanced Access is an opportunity to evaluate and work with Relativity features prior to General Availability release. Relativity customers typically participate in Advanced Access programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments and may not represent the functionality, appearance or behavior of the General Availability release version of this feature.

14.1 Release notes

This section includes the release information and the current functionality of the Search Q&A application.

Release: July 10, 2025

- **Saved Search Indexing**—indexing documents from public saved searches is now supported.

Release: May 23, 2025

- **Multi-turn chat**—simulates human conversation by preserving context from previous exchanges within the same chat session, which enables follow-up questions and deeper exploration without the need for repetition.

Release: April 22, 2025

- **Chat persistence**—chat history is now saved and available for review even after user logs off their Relativity session.
- **Start a new chat**—this button in the Q&A panel starts a new chat conversation below the previous one, separated by a line. The new chat conversation does not reference previous conversations for its responses.

Release: December 9, 2024

- **Search Q&A button**—the button is located on the top of the document list page and available for all users in the workspace to use.
- **"Ask a question" textbox**—when the Q&A panel is open, a user can input a natural language question about the documents. Each question is evaluated and answered individually without consideration of the full conversation.
- **Workspace-level context**—all indexed documents available in the workspace, up to 100,000, are subject to the querying process.
- **Chat history**—a user's chat history is visible during that user's Relativity session. Once the user logs out of Relativity the chat history will be cleared.
- **Citations**—references and citations are provided for each response. A user can filter the Document List to the show these references or can open each reference from the chat panel in the Document Viewer. Please note, that in the initial AA release, we do not confirm that the citations are grounded in the document text like in aiR for Review. See below on future feature releases.

14.2 Considerations

During AA, Search Q&A supports:

- Up to 50,000 documents per index.
- Up to 5 MB of Extracted Text per document.

- Up to 5 workspaces per Advanced Access participant.

14.3 Installation

During AA, the Relativity Search Q&A team will handle the installation. If you are an AA customer, you must provide the following information:

- RelativityOne Tenant
- Written approval for the Relativity team to access the workspace.

After the above items are provided, the Relativity team will install the application. This takes approximately 1-2 business days. Once the application is installed, you will be able to index the data and start using Q&A.

Installation completion details will be emailed to you.

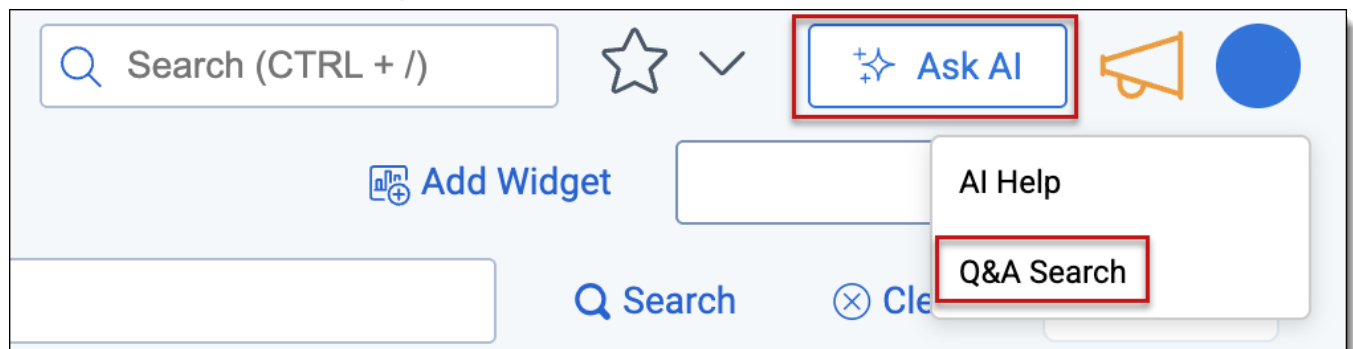
14.4 Prerequisite

During AA, aiR for Case Strategy must be installed and enabled on the workspace for the RelativityOne Tenant to use Search Q&A functionality.

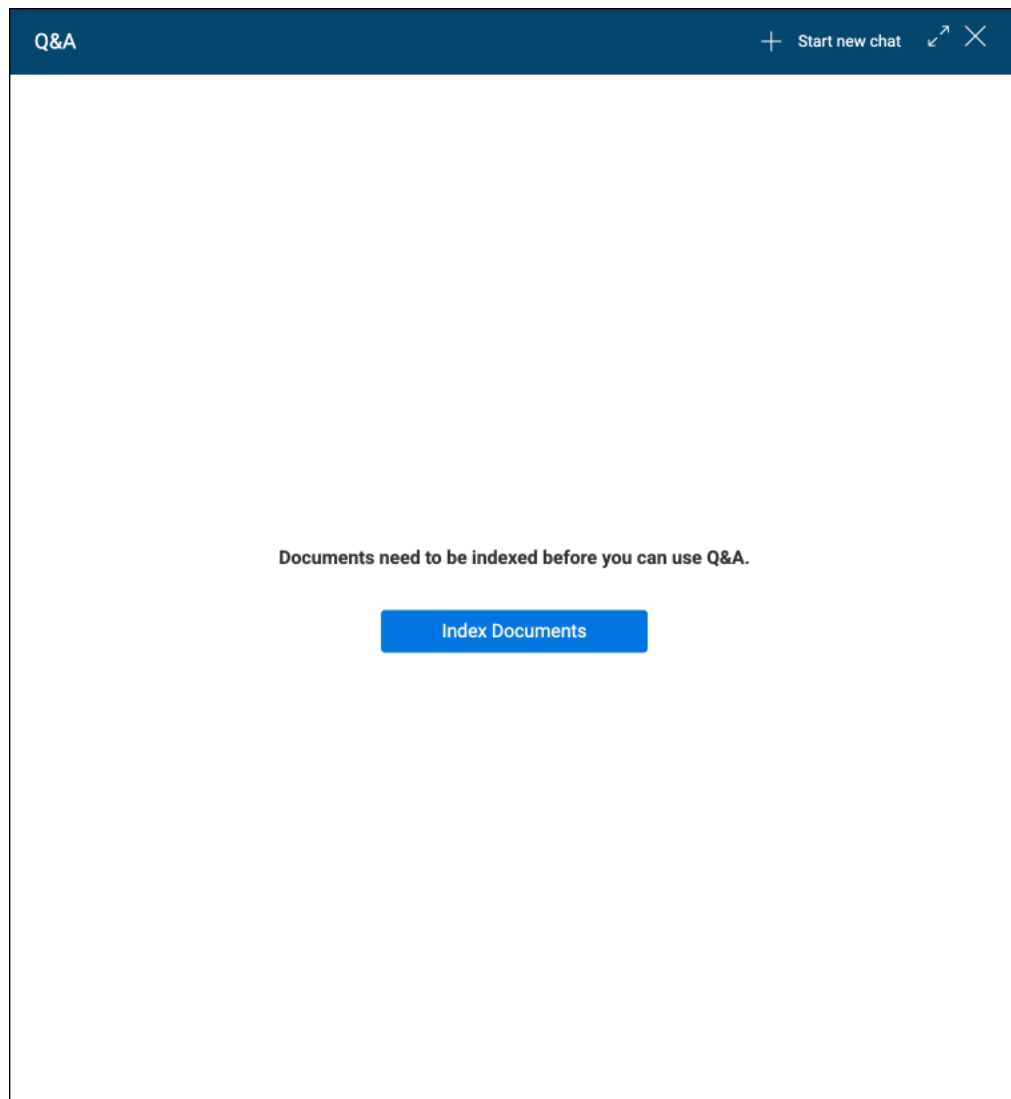
14.5 Using Search Q&A

Follow the steps below to use Search Q&A.

1. Click **Ask AI** at the top of the page, then select **Q&A Search**.



2. Click **Index documents** to begin the indexing process.



3. Select the **Data Source** (public saved search) containing the documents from which to build the index.

Create Q&A Index

Data source ⓘ *

Data Source ▼

Name ⓘ *

Select a data source first

Create index

Cancel

Note: Q&A indexing allows up to 50,000 documents per saved search. If this limit is exceeded, an error message will appear. To continue, decrease the number of documents and restart the indexing process.

4. Enter a descriptive **Name** for the index.
5. Click **Create Index** to start building the index.
The index list displays. This sortable grid contains a detailed overview of the indexes, including build status, number of documents in the index, last update date/time, and name of person who last modified it.
6. Do any of the following on the index list modal:

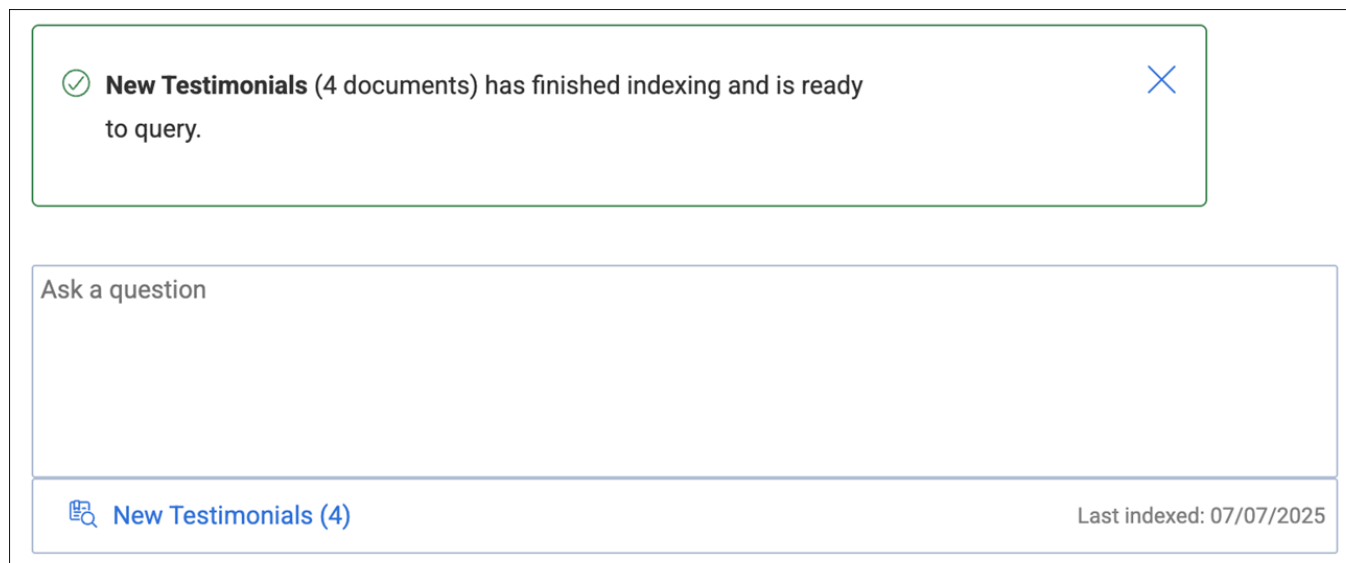
Select Q&A Index

Name ↑	Status	Last Updated	Nr of Documents	Last Updated By
<input checked="" type="radio"/> New Testimonials	Indexed	07/07/2025, 17:14:49	4	Krajewska, Malgorzata
<input type="radio"/> str test	Indexed	04/07/2025, 21:32:54	2,096	Krajewska, Malgorzata
<input type="radio"/> Testimonials	Indexed	06/07/2025, 22:26:46	2,095	Krajewska, Malgorzata

+ New IndexDelete Index

SetCancel

- Select the index to use for querying and click **Set**. Then, proceed to the next step.
 - Click **+New Index** to build another index and repeat the previous steps.
 - Click **Delete Index** to delete one from the list. See [Deleting an index below](#).
7. Once the indexing process completes, enter your question in the *Ask a question* text box using natural language, as if chatting with a person, and click **Enter**. The index in use and when it was last indexed appears below the question text box.




The screenshot shows a chat window. At the top, a green-bordered box contains a checkmark icon and the text: "New Testimonials (4 documents) has finished indexing and is ready to query." with a blue 'X' close button in the top right. Below this is a large text input area with the placeholder "Ask a question". At the bottom of the chat window, there is a bar with a magnifying glass icon and the text "New Testimonials (4)" on the left, and "Last indexed: 07/07/2025" on the right.

8. To use a different index for querying, click the current index name link, then select another one from the index list and click **Set**.

After submitting your question, a natural language response appears that includes a list of references, documents with control numbers, and more. Review the accuracy of the answer and its citations. For additional information on responses, see [Working with responses on page 192](#).

The system preserves the context of your ongoing chat exchange within the current chat session, allowing follow-up questions and simulating human conversation.

To start a brand new chat session, click . A line separates the old conversation from the new one. The new chat does not reference previous sessions to answer new questions.

Chat history remains available even after you log out of the Relativity session.

Note: After each question and response, review the accuracy of the answer and its citations. Where appropriate, record your feedback and send it to the Relativity team.

14.6 Deleting an index

Use the steps below to easily delete indexes from the index list. Deleting an index does not erase past questions and answers, which remain accessible in your chat history.

1. Navigate to the index list from the chat window by clicking the index name link below the text box to open the index list.

2. Select the radio button corresponding to each index to be deleted.

Select Q&A Index

	Name ↑	Status	Last Updated	Nr of Documents	Last Updated By
<input checked="" type="radio"/>	New Testimonials	Indexed	07/07/2025, 17:14:49	4	Krajewska, Malgorzata
<input type="radio"/>	str test	Indexed	04/07/2025, 21:32:54	2,096	Krajewska, Malgorzata
<input type="radio"/>	Testimonials	Indexed	06/07/2025, 22:26:46	2,095	Krajewska, Malgorzata

+ New Index

Delete Index

Set

Cancel

3. Click **Delete Index**.
A warning message appears to let you know the delete action cannot be undone.
4. Click **Delete** to proceed with permanently deleting the selected indexes.

14.7 Q&A panel

Refer to the following table to learn how to use the Search Q&A panel.

Q&A

+

Start new chat

↗

✕

AI generated content may be incorrect, check for accuracy.

4

What Big Thorium is?

5

Big Thorium is a corporation that focuses on providing innovative programs and benefits to its employees, such as engagement initiatives, wellness programs, and professional development opportunities.[1,2]

✓ 2 references

View in the Document List

Testimonials (2095)

6

AI generated content may be incorrect, check for accuracy.

Do you see any financial reports in this dataset?

Yes, the dataset contains references to financial reports.[1,1] The financial reports are mentioned in the context of review and approval for an upcoming board meeting.[1]

> 2 references

View in the Document List

New Testimonials (4)

1 BTV5_000005018

"I have reviewed the updated documents and can confirm their accuracy."

2 BTV5_000005018

"I would like to request your review of the final version of the financial reports that we intend to present at the upcoming board meeting."

Ask a question

7


New Testimonials (4)

8

Last indexed: 08/07/2025

9

Element	Description	Purpose
1	Start new chat	Click to starts a new chat conversation below the previous one, separated by a line. The new chat conversation does not reference previous conversations for its responses.
2	Expand	Click to enlarge the Q&A panel to improve answer readability.


Relativity one

Searching Guide

191

Element	Description	Purpose
3	Close	Click to close the Q&A panel.
4	Question asked	The question asked.
5	Response given	The response to the question asked. See Working with responses below for more information.
6	Index used	The index used for the query and number of documents within it.
7	Ask a question	Type your question into this text box and click Enter to submit it. Questions must be at least three characters long to submit.
8	Current index	The name of the current index being used for the question and number of documents within it. Click the link to change the index used or to create a new index.
9	Last indexed date	The date on which the current index was most recently generated.

14.7.1 Working with responses

After submitting questions, you will receive responses and additional information, as described below


1

Yes, the dataset contains references to financial reports.[1,1] The financial reports are mentioned in the context of review and approval for an upcoming board meeting.[1]


2

> 2 references

3

 View in the Document List

5

 **New Testimonials** (4)

1

BTv5_000005018

4

"I have reviewed the updated documents and can confirm their accuracy."

2

BTv5_000005018

"I would like to request your review of the final version of the financial reports that we intend to present at the upcoming board meeting."

Element	Description	Purpose
1	Natural Language Response	Provides the summarized answer to the question based on the retrieved documents, with inline citations (for example, [1], [2]) to specific References.
2	List of References	Displays the number of referenced citations generated for the response. Click this link to show or hide the list of document references. The first three sentences of the reference displays. Click More to view additional content.

Element	Description	Purpose
3	View in the Document List	Filters the Document List to display only the referenced documents. Note: The selection of a specific folder in the Folder Browser may affect the visibility of references in the Document List. To ensure all references are visible, select the root folder level.
4	Control Number	Click this link to open the document in the Viewer.
5	Index Name	Name of the index used for the query and the number of documents within the index.

14.8 Troubleshooting

Search Q&A may sometimes fail to understand your question or provide a meaningful response. The current version has limited error-handling capabilities. If so, you might see this error message: *I didn't quite understand your question, please try asking it differently.* Try rephrasing the question and asking again. If the issue persists, report it to search-qa-help@relativity.com for troubleshooting.

I didn't quite understand your question, please try asking it differently.

14.9 Removing the Search Q&A application

Currently, the Search Q&A application is available to all users on the workspaces where this capability enabled. Control access to those workspaces with the standard workspace access permissions. To disable Search Q&A from a given workspace, email the Search Q&A team at search-qa-help@relativity.com.

14.10 Frequently asked questions

Listed are questions you might have using the Search Q&A application.

What types of questions can I ask in Search Q&A?

Search Q&A can handle questions related to the documents in the current workspace. Questions that answer general knowledge questions are not supported. You can ask questions in natural language, such as "Who are the key people in this case?" or "Which documents mention [topic]?"

Can Search Q&A understand multiple languages?

Search Q&A primarily works in English. If your workspace contains documents in other languages, the effectiveness of answers may vary.

How is the accuracy of responses ensured?

Search Q&A relies on indexed data to generate answers leveraging retrieval-augmented generation (RAG). While it aims to offer correct responses, always verify the offered information and citations for accuracy. Feedback can help improve the tool over time.

Is there a limit to how many questions I can ask in one session?

There is no set limit, but long or highly complex queries may take more time to process. For best results, keep your questions concise and focused.

Can multiple users access Search Q&A simultaneously?

Yes, multiple users can use Search Q&A within the same workspace concurrently. However, each user's session and chat history are independent.

15 Field categories

Field categories allow you to quickly and accurately apply field conditions to the document list, a saved search, or a view. Once a field category has been created, it can be linked to as many fields as desired. The field category can then be selected when adding fields or conditions in the search panel, a view, or a saved search to conveniently access and apply the linked fields.

All Fields is a system default field category that displays every non-secured field in your workspace. It cannot be edited or secured from other users and it will always display at the bottom of the Field Categories drop-list. All Fields allows you to conveniently switch between viewing all of the fields and the fields included in a particular field category.

Field categories have object permissions so that they can be created and limited to a specific group or workflow within in Relativity. Securing a field category only determines whether it is visible in the drop-list of field categories that can be selected when adding conditions in the search panel.

Note: Field categories that you create are carried over during workspace upgrades and when those cases are used as templates. The All Fields field category will also carry over after an upgrade, though it is currently a system default that cannot be configured by users.

15.1 Creating a field category

To create a new field category, do the following:

1. Navigate to the **Field Categories** tab.

Note: All Fields is a field category that is always present in Relativity but it does not display on the Field Categories tab since it cannot be edited.

2. Click **New Field Category**.
3. Enter the desired information for the following fields:
 - **Name** - the title of the field category
 - **Order** - determines where the field category appears within the view drop-down list. Items that share the same value are sorted in alphanumeric order.
 - Order can be any integer (positive or negative). No decimals are allowed. You can use the **View Order** button to see how items are currently ordered.

Note: Ordering by 10's, starting with 10, then 20, then 30, etc. is recommended because you can insert an item into any position later in the workspace, without the need to reorder all items.

4. Click **Save** to create the new field category.

15.2 Viewing fields associated with a field category

Before linking and unlinking fields to a field category, it can be helpful to see what fields are associated with each field category. To do this, add the Fields field to the desired field category view and select that view on the Field Categories tab.

#	Name	Order	Fields
1	Review Fields	10	Batch Batch::Assigned To Batch::Batch Set Batch::Status Control Number Has Images Has Native
2	Family Fields	20	Family Group MD5 Hash
3	Native and Image Field Category	30	Control Number Extracted Text File Extension Has Images Has Native

All 3 Copy ^ Total: 3

15.3 Link a field or fields to a field category

To link a field or fields to a field category, do the following:

1. Navigate to the **Field Categories** tab.
2. Click on the desired field category in the list.
3. Click **Link**.

Select Items - Fields

1 - 10 of 365 10 per page

#	Name	Field Type	Notes	Key
1	Control Number	Fixed-Length Text	Unique identifier for native files using a document level numbering scheme or first page of a document using a page level numbering scheme.	System
2	Extracted Text	Long Text	Complete text extracted from content of electronic files or OCR data	System

0 - 0 of 0 10 per page

No data.

Apply

Cancel

4. Filter the field list as needed and check any desired fields.
5. Once the desired fields have been selected, click the **Move selected left to right** icon.

- Click **Apply**.
The selected fields are linked to the field category and display in the Fields section.

15.4 Unlink a field or fields from a field category

To unlink a field or fields from a field category, do the following:

- Navigate to the **Field Categories** tab.
- Click on the desired field category in the list.
- Click the checkbox on the left side of the row for any fields you wish to unlink.

Fields							
Link Unlink							
1 - 7 of 7		10 per page					
	Object Type	Name	Field Type	Include in Text Index	Is Relational	Keywords	Notes
	(All)	Filter	(All)			Filter	Filter
<input type="checkbox"/>	Document	Batch	Multiple Object			System	
<input type="checkbox"/>	Document	Batch::Assigned To	User			System	
<input type="checkbox"/>	Document	Batch::Batch Set	Single Object			System	
<input type="checkbox"/>	Document	Batch::Status	Single Choice			System	
<input type="checkbox"/>	Document	Control Number	Fixed-Length Text			System; Relativity Template	Unique identifier for native files using a document level numbering scheme or first page of a document using a page level numbering scheme.
<input type="checkbox"/>	Document	Has Images	Single Choice			System; Relativity Template	
<input type="checkbox"/>	Document	Has Native	Yes/No			System	

- Click **Unlink**.
- Click **Unlink** again in the pop-up.
The selected fields are no longer linked to the field category.

16 Optimized indexing

Optimized indexing requires some knowledge of your data. The time it takes to scrub your data before indexing will be rewarded in time saved when creating an index and returning search results. Consider the following when creating an index:

- Consider removing file types that have no searchable content, such as system or program files.
- Use a separate index for searching database files and large Excel files.
 - Even if your database has only a small number of these files, creating an index without them improves searching speed, especially numeric range searching.
- Set up multiple dtSearch indexes, including one with a smaller document set based on one or more of the following criteria:
 - date ranges
 - custodians
 - text size (extracted or OCR text)
 - Small (< 2 MB)
 - Medium (> 2 MB and < 10 MB)
 - Large (> 10 MB and < 25 MB)
 - Very large (> 25 MB)
- Set a dtSearch index to recognize and/or ignore words, characters, and digits as necessary.
 - Noise Words (Ex: Include “sample” as a part of the noise word list for a dtSearch index containing Excel documents.)
 - Alphabet file list (Ex: Index the character “£.”)
 - Remove numbers from the alphabet file list if only searching for words – this reduces the size of the index and disables numeric range searching.
- Enable dtSearch indexes to automatically recognize dates, email addresses, and credit card numbers only when necessary. Enabling this setting increases build time.
- Consider using a pair of dtSearch indexes when adding new data. You can have one index updated in the background and then swap out the outdated index with the current one.

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.

©2025. Relativity ODA LLC. All rights reserved. Relativity® is a registered trademark of Relativity ODA LLC.