



aiR Assist Guide

October 7, 2026

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

Table of Contents

1 Assist (aiR Assist)	4
1.1 Assist availability	4
1.2 Index and document limits	5
1.3 Basic getting started steps	5
1.4 How Assist finds answers	6
1.5 Understanding Assist responses	7
1.6 Language support	7
1.7 Common use cases	8
1.8 Auditing user activity	8
2 Permissions	10
2.1 Considerations	10
2.2 Typical permissions scenarios	10
2.3 Setting permissions	11
3 Best practices	12
3.1 Prompting	12
3.1.1 Examples of unsupported prompts and their supported alternatives	13
3.2 Indexing	14
4 Prompting	15
4.1 Considerations	15
4.2 Navigating the prompt composer	15
4.3 Navigating the responses	16
4.4 Submitting prompts	18
4.5 Starting a new prompt conversation	19
4.6 Starring (bookmarking) a prompt and response	20
4.7 Copying a response	20
4.8 Conversation history and sessions	20
5 Intent detection	20
6 Feedback options	22
6.1 Feedback controls for responses (thumbs up/thumbs down)	22
6.2 In-product feedback survey	22
7 Using Conversation Manager	22
7.1 Navigating the Conversations panel	23

7.2 Opening conversation manager	23
7.3 Viewing a conversation	24
7.4 Renaming a conversation	24
7.5 Deleting a conversation	24
7.6 Pinning a conversation	25
7.7 Exporting a conversation	25
8 Management	25
8.1 Indexes created in Case Strategy	26
8.2 Indexes created within Assist	26
8.3 Throughput expectations	26
9 Working with indexes	27
9.1 Navigating the Index Manager	27
9.2 Indexing rules and limitations	28
9.3 Creating a new index	28
9.4 Rebuilding an index	30
9.5 Editing an index name	31
9.6 Deleting an index	31
9.7 Monitoring the index build process	32
9.8 Error handling and retrying index jobs	33
10 Metadata mapping	34
10.1 Considerations	35
10.2 Configuring metadata mapping	35
10.3 Using metadata for prompting	36
11 Frequently asked questions	36
11.1 Prompt FAQs	36
11.2 Response FAQs	37
11.3 Conversation FAQs	38
11.4 Index FAQs	39
11.5 Document FAQs	39
12 Index	41

1 Assist (aiR Assist)

Note: aiR Assist is now Assist. References to aiR Assist may still appear in this documentation, in the product, and in other sources such as marketing materials, published content, and commercial or legal agreements. Both names refer to the same product/feature.

Assist is a conversational search tool integrated within a Relativity aiR workspace, designed to empower legal teams to interact with their data using natural language. By leveraging advanced AI, Assist helps to surface potential insights, reveal possible connections, and uncover themes more efficiently. This can enhance the process they use to analyze and interpret legal data more effectively, potentially leading to quicker understanding, better decisions, and defensible workflows when validated by users.

It works by searching the extracted text of indexed documents. Users can create up to five indexes per workspace, each supporting up to 300,000 documents. When a query is submitted, Assist identifies the documents deemed most relevant and employs a large language model (LLM) to generate answers grounded by citations tied back to the documents.

Note: ARM is not supported for Assist. Assist permissions, indexes, conversations, and metadata mapping cannot be archived, restored, or moved using ARM.

For additional information on AI and aiR products, see AI products and features in Relativity aiR on the help site.

Refer to [Frequently asked questions on page 36](#) for helpful answers to common questions about using Assist.

1.1 Assist availability

Assist is available to organizations with an Case Strategy contract or aiR Integrated contract. It does not require installation in Relativity aiR.

Contract type	Availability
Case Strategy contract	Assist is included in Case Strategy at no additional cost and does not draw down on document thresholds. It is available in review workspaces where Case Strategy is installed and at least one Fact Extraction or Transcript Summary job has completed. Customers can also create a Case Home index. For more information about completing these jobs, refer to Case Strategy documentation. Note: If an Case Strategy workspace does not meet the required conditions (Case Strategy is installed and a Fact Extraction or Transcript Summary job has completed), users can open Assist, but they cannot use it until one of those jobs runs.
aiR Integrated contract	Assist is included in Integrated contracts at no additional cost and does not draw down on document thresholds. It is available in both repository and review workspaces. The Case Strategy installation and job-completion requirements do not apply. Customers can also build a Case Home index based on Case Strategy documents. For more information about completing these jobs, refer to Case Strategy documentation.

To use Assist, the appropriate permissions must be configured for each group in each workspace, regardless of contract type. See [Permissions](#) for more information. By default, System Administrators are granted access to *all* Assist permissions across *all* workspaces.

The Assist icon appears at the top of the sidebar.

Note: If you have an integrated pricing contract, you can access Case Strategy and Assist in both standard and repository workspaces. For Case Strategy, the integrated pricing contract provides an allotment of up to 50,000 documents to use in each repository workspaces for each period, generally annually, of the SaaS subscription term. If a user tries to start a job that has enough documents to put you over the 50,000 document limit, the job fails to run. For aiR Assist, there are no additional considerations for repository workspace use on integrated pricing contracts. To learn more, contact your Account Executive.

1.2 Index and document limits

Below are the index and document limits for Assist.

Type	Limit
Index	- Each index can contain up to 300,000 documents. - A maximum of five (5) built indexes can exist at one time per workspace. This can be from a combination of the Case Home document set (created in Case Strategy) and public saved searches. See Working with indexes on page 27 for details on adding and deleting indexes.
Document	- Individual documents must be up to 5 MB of extracted text; larger files are excluded during indexing. - Only documents with extracted text are indexed. The text must be stored in Data Grid (not SQL). Documents that do not have extracted text are automatically excluded from the index.

1.3 Basic getting started steps


Some initial configuration must be completed before anyone can begin using Assist. Below are some basic workflow steps to getting started with and using Assist.

Task	Description
1. Configure user group permissions	User(s) granted Access Management permission (typically a System Admin) assigns permissions to users, such as Index Management and Prompting. See Permissions on page 10 for details.
2. Map the metadata	If using metadata, user(s) granted Index Management permission (typically a System Admin) maps the metadata fields. See Metadata mapping on page 34 for details.
3. Create indexes	User(s) granted Index Management permission creates

Task	Description
	1-5 indexes (document sets) based off of public save searches that will be used by those with Prompting permission to generate responses to their prompts. Working with indexes on page 27 .
4. Run a Fact Extraction or Transcript Summary job (only for Case Strategy contracts)	If your organization has an Case Strategy contract, at least one Fact Extraction or Transcript Summary job must be completed. See Case Strategy documentation for details.
5. Ask questions/prompts	Users granted Prompting permission enter questions/prompts against a selected index to generate responses and document citations. See Prompting on page 15 for details. Also refer to Best practices on page 12 for guidance on creating prompts and Using Conversation Manager on page 22 for managing conversations.
6. Review and verify responses	Users granted Prompting permissions review and verify the responses in the Conversations panel, Document list, and Viewer. See Navigating the responses on page 16 for details. Also refer to Best practices on page 12 for guidance on indexing.

1.4 How Assist finds answers

Assist operates using a Retrieval-Augmented Generation (RAG) process to deliver grounded, evidence-based responses. This approach combines document retrieval with LLM generation to help support accuracy, transparency, and contextual relevance.



- Asking a question (question step)**—The user asks a question aimed at the index (document set) they selected to use.
- AI analyzes the question and plans the search (analyzing step)**—The question is read by the LLM, which decides whether the question is an appropriate one to answer. If so, it formulates queries to run against the index to find documents.
- Searching relevant documents (search step)**—Assist performs one or more searches to identify the most relevant content within the index of document chunks.
- Generating the answer with links to cited documents (generation step)**—The selected passages, along with the original question and system prompt, are passed to the LLM. The model uses this retrieved context to generate a response intended to be coherent, concise, and supported by retrieved content, citations, and references to the original sources.
- Links to cited documents (answers step)**—The user reviews and verifies citations in linked documents.

Note: Assist responses are based only on the documents the requesting user is allowed to access at the time they ask the question. Because document-level access can vary by user, different users may receive different answers to the same question.

1.5 Understanding Assist responses

Assist helps identify and summarize potentially relevant information across large document sets using natural language interaction. Built on a Retrieval-Augmented Generation (RAG) architecture, it retrieves and analyzes the documents most likely to be relevant, then generates a citation-supported response based on that content.

It returns contextually relevant and evidence-based information rather than performing exhaustive or “find everything” searches. Because it does not review each document individually, some keyword or topic matches may not be included in the response.

The RAG process works best when key evidence is found in a few focused documents. Results are less accurate if answers depend on scattered or unclear information.

It is also important to note that document security may differ between users. As a result, two users could ask the exact same question and receive different answers, because the LLM interaction used for each response only knows about documents that the requesting user is authorized to access. Assist does not retain knowledge from other users' conversations, and every interaction is evaluated against the requesting user's document security permissions at the time of the prompt.

1.6 Language support

The LLM used by Assist has been evaluated for use with 83 languages. Although Assist has primarily been tested on English-language documents, it is designed to support non-English datasets. For more information, see [Language support in aiR products on the Relativity help site](#).

Unlike other aiR products, Assist system consists of two components:

- **Retrieval**—the LLM generates search queries to find relevant documents in the index.
- **Generation**—the LLM summarizes the retrieved documents, picks relevant ones, and synthesizes them into an answer.

Because of this two-step design, there are special considerations when working with non-English datasets. If you use Assist with non-English datasets, we recommend the following:

- When possible, write your question in the same language as the documents being queried—ideally one you are fluent in. If that is not possible, write your question in English.
- By default, Assist searches for documents in English. If your documents are in another language and you are writing your question in English, tell it which language to search in by adding a sentence such as: "Search in <Language>." This helps the retrieval step generate queries in the specified language. While effective in most cases, the LLM may not always apply the requested language consistently.
- Extracted text and citations remain in the language of the source document — you do not need to translate them yourself before reviewing.
- By default, answers are generated in English. To request an answer in another language, add a sentence to your question, such as "Write answer in <Language>." This approach is effective in most cases, but the LLM may not always generate answer consistently in the requested language.

You can inspect which searches were attempted using the Search Completed summary in the response panel (see [Navigating the responses on page 16](#)). Use this to iterate on your question and adjust language coverage as needed.

1.7 Common use cases

Here are some example questions targeting a few common use cases for Assist:

Use case	Common category	Example question
Early Case Insight	Finding potentially key/important documents	<i>Can you find me documents that discuss potential gifts or incentives?</i>
	Finding documents by theme	<i>Are there any documents mentioning fraudulent behavior of John Doe?</i>
	Understanding actors and roles	<i>Who was involved in discussions about offering gifts?</i>
Case Strategy Development	Identifying a series of events	<i>Create a high-level timeline for events that took place before the start of Project Artemis.</i>
	Understanding communications and relationships between actors	<i>Who communicated with whom about the contract terms?</i>
Deposition/Trial Preparation	Suggesting exhibits based on key criteria	<i>List documents to use as exhibits based on [key document criteria].</i>
	Confirming conversations or actions took place	<i>Did John Maxwell send an email about the compliance policy?</i>

Additional scenarios include but are not limited to:

- **Evidence identification**—query an issue to quickly surface relevant, citation-backed documents.
- **Communication analysis**—map interactions and relationships between custodians and actors.
- **Production deficiency review**—test incoming productions against RFP requests to identify gaps.
- **Entity and role understanding**—clarify who's who and their responsibilities across the document set.
- **Sample set identification**—find candidate documents to use in the relevant sample set for an aiR Review project.

1.8 Auditing user activity

You can monitor Assist user activity in the Audit application in your Relativity instance. The Audit table records the events listed below. Because Assist audit events are not RDO objects, they use an **Audit Object UUID** in the Audit table instead of an Object Artifact ID. For more information on using the Audit application, see [Audit documentation](#).

Note: This is a temporary solution, therefore, we do not recommend building any integrations based on these audit record types. The data is planned to be moved to Custom Reports by the end of 2026.

Click to see the Assist object types and corresponding actions that will appear in the Audit table:
When you filter on the Object Type column in the Audit application, RDO objects are listed first. Since Assist is not an RDO object, its object types appear at the end of the list. To locate them faster, enter *Assist* in the filter search box.

Note: Due to the September 2026 product rebrand, Object Type names appear in two formats depending on whether they were used before or after the September 17, 2026 changeover date. Use the columns below to find the correct names.

Object Type name <u>prior to</u> September 17, 2026	Object Type name <u>after</u> September 17, 2026	Action
aiR Assist Question	Assist Question	• Create—records auditing data when a question (prompt) was submitted from the Ask a question box. Audit does not record the content of the question.
aiR Assist Answer	Assist Answer	• Create—records auditing data when the answer to a question (prompt) was generated. Audit does not record the content of the answer.
aiR Assist Conversation	Assist Conversation	• Create—records auditing data when a new question and answer conversation started (as in, when the user clicks the New conversation icon). • Delete—records auditing data when the conversation was deleted. • Edit—records auditing data when the conversation was edited (renamed). • View—records auditing data when the conversation was opened for viewing from the Conversations panel. • Export—records auditing data when the conversation was exported.
aiR Assist Index	Assist Index	• Create—records auditing data when the index was created. • Delete—records auditing data when the index was deleted. • Update—records auditing data when the index was updated (rebuilt) or its name was edited.
aiR Assist Indexing Error List	Assist Indexing Error List	• View—records auditing data when the indexing error page was viewed.
aiR Assist Metadata Mapping	Assist Metadata Mapping	• Update—records auditing data when the metadata mapping information was edited.

So, for example, if you opened a conversation, submitted a new question (prompt), and received an answer (response), the records would be as follows:

- Assist Conversation = View
- Assist Question = Create
- Assist Answer = Create

2 Permissions

System administrators can control access to Assist per group within a workspace using Features permissions. This is helpful when different teams are responsible for different tasks, such as only using indexing or only asking questions. See Feature permissions and Setting workspace permissions for more information.

Note: ARM is not supported for Assist. Assist permissions, indexes, conversations, and metadata mapping cannot be archived, restored, or moved using ARM.

2.1 Considerations

- By default, users in the System Administrator group have access to *all* Assist permissions across *all* workspaces. These permissions *cannot* be modified. Users in the Workspace Administrator group will not automatically see Assist unless their permissions are configured, as described below, or their permissions were included in the template used for the workspace.
- Feature permissions can be included in a template and copied to new workspaces created from that template. It does not retroactively set permissions on existing workspaces.
- If Access Management is the *only* permission selected for a group, Assist will *not* appear in the workspace. At least one of the other permissions must be enabled for Assist to appear.
- Since document security may differ between users, Assist responses to prompts are based on the documents in the index (data source) that the requesting user is permitted to access at the time of the prompt submission. See [Understanding Assist responses on page 7](#) for more details.

2.2 Typical permissions scenarios

Click to see some typical scenarios, along with the appropriate permission settings:


Group	Permitted Capabilities	Index Management Permission	Prompting Permission	Conversation Export Permission	Access Management Permission	Can see Assist?
Group A	Full Assist usage: • Index management • Edit metadata mapping • Prompt generation • Conversation management	Yes	Yes	Yes*	No	Yes

Group	Permitted Capabilities	Index Management Permission	Prompting Permission	Conversation Export Permission	Access Management Permission	Can see Assist?
	and export					
Group B	- Index management (create, rebuild, delete) - Edit metadata mapping	Yes	No	No	No	Yes
Group C	- Prompt generation - Conversation management and export	No	Yes	Yes*	No	Yes
Group D	- Only manage permissions for other groups - No access to Assist functionality - System Admin level	No	No	No	Yes	No

* - Users can only export the conversations they created using Prompting.

2.3 Setting permissions

Note: This feature supports feature permission actions. Each action grants access to a specific task within the feature. You can grant groups access to individual actions, allowing them to perform only the tasks they need. To grant access to the entire feature, select all of its actions. Manage feature actions on the Features tab in a group's workspace permissions. For instructions, see [Setting workspace permissions](#). For more information about feature permissions, see [Feature permissions](#).



Permissions for Assist include:

Note: Assist will not appear in the group's workspace if Access Management is the *only* permission selected for the group. At least one of the other permissions must be enabled for Assist to appear.

- **Access Management**—select to allow group members the ability to assign and manage Assist permissions for other groups in the workspace. If this is the only permission enabled for the group, Assist does not appear for them.

Caution: Because this advanced setting affects group permission assignments, System Administrators should enable it only for select groups.

- **Index Management**—select to allow group members to use Index Manager to create, rebuild, and delete indexes and map metadata in the workspace. If this is the only permission enabled for the group, users cannot submit prompts, and they cannot start and manage conversations. For more information on Index Manager, see [Management on page 25](#).
- **Prompting**—select to allow group members to submit prompts about documents in the workspace using created indexes and optionally use metadata for results. They can also manage their conversations through the Conversation Manager. For more information on prompting and conversation management, see [Prompting on page 15](#) and [Using Conversation Manager on page 22](#).
- **Conversation Export**—select to allow group members to export their own conversation content to a Word (.docx) file. See [Exporting a conversation on page 25](#) for more information.

3 Best practices

Refer to the prompting and indexing best practices below to effectively use Assist.

3.1 Prompting

Following these prompting best practices may help ensure accurate, efficient, and relevant responses from Assist.

- **Ask clear and focused questions**
To help improve results, keep questions concise and specific. Focus each query on a single topic or piece of information so it can be mapped to documents. Long, compound, or highly complex questions may take longer to process, may not map to documents, and can reduce clarity in the generated response. Start a new conversation when switching topics to prevent stale content from skewing results.
- **Leverage keywords and synonyms**
Retrieval engines like specific terms. Include specific names, likely variations of keywords (such as, bribe, gift,

incentive, Rolex), or entity synonyms (such as, bt or bt.us for Big Thorium). Retrieval benefits from alternative phrasing.

- **Stay within the saved search context**

Assist generates responses based on the documents included in the public saved searches used to build its indexes. Each saved search defines the specific dataset Assist can draw from within the workspace. Questions should therefore relate to the content of those saved searches rather than general or external topics.

- **Expect some variation in repeated queries**

Submitting the same or similar questions multiple times may produce slightly different answers, as Assist regenerates responses dynamically. However, the core content and conclusions are generally expected to remain similar, though variation may occur.

- **Review citations and supporting references**

Each Assist response includes citations and supporting document references. Review these sources to verify accuracy and context, especially when using the results for analysis, reporting, or decision-making.

- **Maintain high-quality indexed data**

Response quality depends on the content indexed. Ensure that the dataset includes clean, text-extractable documents and that saved searches accurately capture relevant materials. Avoid including duplicate or irrelevant documents within indexes.

- **Avoid overly broad or “find me everything” queries**

Assist is optimized to find and synthesize the most relevant information, not to return exhaustive lists of *all* matching documents. For comprehensive discovery, you can use standard search tools in combination with Assist or, depending on the use case, consider using our other aiR Suite products.

- **Avoid using Assist for calculations**

Do not rely on Assist to add up invoices, total damages, or perform complex calculations. The right data may be scattered across multiple files, and LLMs are not perfectly reliable at math calculations.

- **Break down complex reasoning or multi-step queries**

Avoid questions that require multiple steps or “leaps” (such as, “show me emails from the director who signed the compliance policy”). Break these into smaller, sequential questions so Assist retrieves the most relevant information and clearly understands your objective.

3.1.1 Examples of unsupported prompts and their supported alternatives

Below are a few examples of unsupported prompts and their supported alternatives:

Unsupported prompt wording example	Why it is unsupported	Supported alternative wording example
Find me all examples of [issue].	Assist is optimized to identify and synthesize the most relevant information, rather than provide exhaustive lists of every matching document. Because it uses a Retrieval-Augmented Generation (RAG) architecture, it retrieves a relevant subset of the corpus and therefore cannot guarantee complete coverage. See Understanding Assist responses on page 7 for more information.	<i>What evidence is there of [issue]?</i> Ask for the evidence or key examples, rather than an exhaustive list.
Find me all conversations by [actor].		<i>Did [actor] discuss [topic]?</i> Focus on a particular topic or issue, instead of asking for all conversations.
Find me all conversations that took place between [actor] and [actor].		<i>What topics did [actor] discuss with [actor]?</i> Focus on the substance or topic of the email exchanges, instead of

Unsupported prompt wording example	Why it is unsupported	Supported alternative wording example
		asking for all conversations.
Find me emails between March 1-March 31, 2012.		<i>What was discussed about [topic] in March 2012?</i> Tie the date to a topic so it becomes a document corpus related question.
Find me emails sent after-hours.		<i>Did [actor] discuss [topic] outside business hours?</i> Anchor it to a topic or actor, instead of an exhaustive list.
Find me the earliest email referencing [term].	Assist retrieves a set of documents and answers questions based on the most relevant results it finds, but that set may not include the earliest or latest email.	<i>What evidence is there of [term] being discussed?</i> Ask about the term to pinpoint the earliest instance.
Write me a prompt for aiR Review related to this matter.	This prompt is not asking a question about specific indexed documents. It is asking Assist to generate content for a task other than document search.	<i>What are the key issues and actors in this matter?</i> Use Assist to summarize the facts, then craft the prompt yourself from that summary.
What is [document name] about?	Assist cannot reference or look up documents by name, hyperlink, or document ID. It only searches extracted text and metadata (emails and date), so a query related to a specific document name does not map to how it retrieves.	<i>What does the corpus say about [subject/keyword from that document]?</i> Query the content or key terms rather than the filename or ID.
What are the discrepancies between [doc 1] and [doc 2].		<i>What did [actor] say about [topic]?</i> asked once per source. Ask separate questions. Break the main question into smaller, sequential questions that can be compared.

3.2 Indexing

To achieve the best indexing results, below are some data preparation recommendations:

- **Build focused indexes**—align each index to specific custodians, issues, themes, or time periods, for example.
- **Curate the Saved Search**—point the index at a well-defined public saved search; private ones are not supported.
- **Optimize document quality**—ensure documents contain extracted text and that the Extracted Text field is Data Grid File System enabled, as Assist does not support environments where the text is stored in SQL. Also, exclude files over 5MB, duplicate documents, irrelevant documents, and documents that do not contain text.
- **Manage indexes well**—establish clear naming conventions, use names longer than five characters, and use the index Description field to help clearly describe the index.

Note: ARM is not supported for Assist indexes.

4 Prompting

You can ask Assist about indexed documents in natural, conversational language. Follow-up prompts build on earlier ones, and Assist uses the conversation context and any people or events mentioned.

Users must be granted Prompting permission to ask questions and manage their conversations. See [Permissions](#) for details on the Prompting permission.

Note: ARM is not supported for Assist. Assist permissions, indexes, conversations, and metadata mapping cannot be archived, restored, or moved using ARM.

Suggested resources:


- [Best practices](#) for prompt-writing tips.
- [Frequently asked questions on page 36](#) for helpful answers to common questions about prompts and responses.

4.1 Considerations

- Prompts and responses are private to each user and persist across the user's sessions. Additionally, they cannot be shared. See [Using Conversation Manager](#) more information on managing conversations. However, you can export the contents of a conversation to an MS Word (.docx) file and share them that way. See [Exporting a conversation on page 25](#) for details.
- It is important to note that document security may differ between users. As a result, two users could ask the exact same question and receive different answers because the LLM interaction used for each response only considers the documents that the requesting user is authorized to access. Assist does not retain knowledge from other users' conversations, and every interaction is evaluated against the requesting user's document security permissions at the time of the prompt. See [Understanding Assist responses on page 7](#) for more details.
- For OCR documents, Assist citations from Extracted Text may not exactly match the source text or display citation highlights due to OCR quality issues, such as extra spaces or misrecognized characters. These OCR-related issues can affect matching, and the LLM may correct minor OCR errors when citing.
- When Relativity updates the underlying large language model (LLM), results may differ even when using the same prompt criteria text because they are tied to the model used when the result was generated. Results generated before a model update use the previous version, while those generated after the update use the new model. Previous results are not automatically reprocessed when LLM models are updated. For additional information, see AI products and features in Relativity aiR on the help site.
- Indexing and prompting do not incur additional costs.

4.2 Navigating the prompt composer

Refer to the following descriptions below to learn how to use the prompt composer (Ask a question box).



The screenshot shows a text input field with the placeholder text "Ask a question about the data in your index". Below the input field, there are two labels: "Progress" with a small icon and "Use Metadata" with a toggle switch that is currently turned on. A small circular icon is visible on the right side of the input field.

- **Ask a question**—type your prompt into this box. Prompts must be at least three characters long to submit.
- **Index used**—name of the index used for the query appears in the lower left corner of the box.
 - Hover over the name to view the saved search source used, number of documents within the index, date/time last indexed and by whom, and description entered when the index was created.
 - Click the index name to display the Index Manager pane where you can select to use a different index for the prompt. See [Working with indexes on page 27](#) for more information on indexes.
- **Use Metadata** toggle—if metadata mapping is enabled in your workspace, the **Use Metadata** toggle displays next to the Index Name. Toggle it on to use metadata for the query or toggle it off to not use it. See [Metadata mapping on page 34](#) for more information.
- **Send message**—after typing your question, click the circled up arrow to submit it for a response.

4.3 Navigating the responses

Refer to the following descriptions below to learn how to navigate the responses generated by Assist.

The screenshot shows the Assist interface with a search prompt and a detailed response. Numbered callouts highlight specific UI elements:

- 1**: The search prompt input box containing the text "which emails mention emergency board meeting".
- 2**: The "Search Completed" status indicator with an expandable arrow.
- 3**: The first item in the search results list, titled "Explicit 'Emergency Board Meeting' email".
- 4**: A reference link (e.g., [1], [2]) within the search results text.
- 5**: The "View in Document List" button located below the search results.
- 6**: The "All Documents" button located at the bottom right of the search results section.
- 7**: The "Ask a question about the data in your index" input box at the bottom of the interface.

At the bottom of the interface, a disclaimer states: "AI-generated content may be incorrect".

1. **Prompt asked**—the prompt that was asked displays above the response.
2. **Search Completed**—click the expand arrow to view a summary that outlines the steps the AI agent took to answer the question.

Search Completed

- Context [All Docs](#)
- Planned approach to answer your question "which email mentioned nuclear power plant construction"
- Found [101 documents](#) - 3 searches
 - Search 1
keywords: nuclear power plant construction regulatory licensing requirements email
 - Search 2
keywords: "construction of the nuclear power plant" email
 - Search 3
keywords: upcoming nuclear power plant project email
- Selected [47 documents](#) after reranking and filtering
- Cited [13 documents](#) in the answer
- Verified that citations are grounded in the documents

- It includes the index name (data source) used, the number of documents found, the searches it attempted, the number of documents selected after re-ranking and filtering, the number of documents cited in the answer, and whether it verified that the citations were grounded in the documents.
- Click the links where available to view specific documents corresponding to that step.

3. **Natural language response and citations**—the responses to the question based on the retrieved documents, with inline numbered citations to specific document references.

- Hover over a citation number (for example 1, 2, 3) to quickly see a preview fragment of the citation from the document. You can click the Document Control Number in the preview to open the document in the Viewer. See Viewer documentation for more information on using the Viewer.
- Click the citation number to open the cited content, highlighted in light blue, on the document's Native or Extracted Text tab in the Viewer. See Viewer documentation for more information on using the Viewer.

Note: For OCR documents, Assist citations from Extracted Text may not exactly match the source text or display citation highlights due to OCR quality issues, such as extra spaces or misinterpreted characters. These OCR-related issues can affect matching, and the LLM may correct minor OCR errors when citing.

4. **# references**—click the expand arrow next to the number of references to show or hide the full list of cited referenced documents generated for the response.

8 references View in Document List

1	EE0000000005	From : Jeff Gilles [jeff.gilles@relativitysandbox.com] Sent : 2/15/2021 8:51:31 PM To : Maks Babuder...
2	EE0000000004	From : Maks Babuder [maks.babuder@relativitysandbox.com] Sent : 2/15/2021 9:30:55 PM To : Jeff Gilles...
3	EE0000000012	Bitcoin \$50K? With Tesla now on board, crypto bulls say the sky's the limit February 9, 2021 Good morning, European stocks and U.S...
4	EE0000000012	ng U.S. ESG ETFs and funds. Will true climate hawks push them to boot out a Bitcoin-endorsing Tesla? UBS Chief economist Paul...
5	EE0000000012	is considered an "intangible asset." If the value declines, that could saddle Tesla with big losses. Will Apple be next? The question...
6	EE0000000040.0001	mat 1 Paid rWnana l ropna, sae,ulme to Check it. Preparer WALTER H DEVHLE, 31d ralyDay ax Use Only CPA Se9-arootoo Pra...
7	EE0000000040.0001	Form 9960 (2019) 17:3 BIDEN JR., JOSEPH JOSEPH R. BIDEN JR. & JILL T. BIDEN FORM 1040 WAGES RECEIVED AND TAXES...
8	EE0000000031	From : Google Payments [payments-noreply@google.com] Sent : 2/15/2021 7:04:21 PM To : maks.babuder@relativitysandbox.co...

- Below the **Document Control Number** link are the first few pieces of information associated with the reference. Click the down arrow associated with each document to view more information.

- Click the **Document Control Number** link to open the document in the Native Viewer or Extracted Text Viewer with the cited content highlighted in light blue. See Viewer documentation for more information on using the Viewer.
5. **View in Document List**—click to filter the Document List to only show documents that include a reference.
 - Be aware that choosing a specific folder in the Folder Browser might change which references appear in the Document List. To ensure all references are visible, select the root folder level.
 - From the Document list, you can create a saved search by using the Save Search button. Refer to documentation on saving searches on the Documents tab for the details.
 6. **Index name**—the name of the index used for the query displays below the list of references. Hover over the name to view the number of documents within the index, date/time last indexed and by whom, and the description entered when the index was created.
 7. **Action buttons**—click the action buttons to do the following:
 - **Star icon**—click to bookmark an important prompt and response pair for quick access later from the conversation header in the Assist panel or from Conversation Manager. See [Starring \(bookmarking\) a prompt and response on page 20](#) for more information.
 - **Copy icon**—click to copy the entire response displayed to the clipboard, including citations. See [Copying a response on page 20](#) for more information.
 - **Thumbs up icon**—click to indicate a positive reaction to the answer provided. See [Feedback controls for responses \(thumbs up/thumbs down\) on page 22](#) for more information.
 - **Thumbs down icon**—click to indicate a negative reaction to the answer provided. See [Feedback controls for responses \(thumbs up/thumbs down\) on page 22](#) for more information.


See [Understanding Assist responses on page 7](#) for more details on how the system operates on a Retrieval-Augmented Generation (RAG) architecture to retrieve and analyze documents for responses.

4.4 Submitting prompts

To submit a prompt (question):

Note: Assist responses are based only on the documents the requesting user is allowed to access at the time they ask the question. Because document-level access can vary by user, different users may receive different answers to the same question.

1. Click the **Assist** icon in the side bar to open the Assist chat panel.
2. Type your prompt in the **Ask a question** box using natural language, as if chatting with a person. Assist formats responses for you, but you can include instructions to use headers, tables, bullets, or numbered lists, which is helpful for longer outputs.



3. Use the default index that displays in the Ask a Question box or click the index name to select a different one from the Index Manager list and click **Set index**.
4. If available, enable the **Use Metadata** toggle to leverage metadata for the prompt or disable it to not use metadata. See [Metadata mapping on page 34](#) for more information.

Notes:

- Users with Prompting permissions enabled can use this toggle. See [Permissions on page 10](#) for more information.
 - The Use Metadata toggle is only supported for indexes created using Assist. Currently, Case Strategy Case Home index does not support metadata mapping so the toggle is disabled.
-

5. Click the **Send Message** up arrow icon to submit the prompt.
Assist displays a response that includes a list of references and documents with control numbers. For additional information on responses, see [Navigating the responses on page 16](#) and [Understanding Assist responses on page 7](#).
 6. Click **View in Documents List** to filter the Document List to only show documents that include a reference. Be aware that choosing a specific folder in the Folder Browser of the Document List might change which references appear in the list. To ensure all references are visible, select the root folder level. From the Document list, you can create a saved search by using the Save Search button. Refer to documentation on saving searches on the Documents tab for the details.
 7. Continue prompting the system using the selected index. If you want to use a different index for querying, click the current index name link, select another one from the index list, and click **Set**. Then, repeat these steps.
-

Note: Responses and citations should be reviewed and verified before using them. Where appropriate, record your feedback using the thumb up or down icon.

The system maintains chat context within the current sessions for follow-up questions and more natural conversation. For more information on managing conversations, see [Using Conversation Manager](#).

Chat history for all chat sessions:


- are private to each user
 - cannot be shared
 - persists across each user's Relativity session
-

Note: When Relativity updates the underlying large language model (LLM), the same prompt criteria text may produce different results because they are tied to the model used when the result was generated. Previous results are not automatically reprocessed when LLM models are updated.

4.5 Starting a new prompt conversation

You can start a new conversation using the same or different index. The new conversation does not reference previous conversations for its responses.

In the Assist chat panel, click the **New conversation** pencil and paper icon and begin typing your query in the box.




For more information on managing individual conversations, see [Using Conversation Manager](#).

4.6 Starring (bookmarking) a prompt and response

You can star (bookmark) important prompt and response pairs to quickly find them later from the conversation header in the Assist panel or from Conversation Manager. For example, you might star key findings or deposition leads to keep them easy to locate among other conversations.

1. Click the **Star** icon below a response to bookmark the prompt and response pair as important. A solid star icon indicates the information has been starred. Click it again to remove the selection.

A **Starred #** badge appears on the conversation banner in both the conversation header and the Conversation Manager, showing how many responses have been starred. The number indicates how many prompt and response pairs are starred in that conversation.



2. Click the **Starred #** badge in either location to display a pop-up list of all starred pairs for that conversation, sorted by most recently starred. Click an **entry** in the pop-up list to immediately go to that prompt and response pair. Click the **Starred #** badge again to close the pop-up list.

4.7 Copying a response

The Copy action lets you copy the full response, including citations, to your clipboard with formatting intact for use elsewhere.

4.8 Conversation history and sessions

Assist maintains a continuous conversation history for each user in the workspace to keep responses consistent. Users can create unlimited conversations, and each user's history remains private and cannot be shared with or accessed by others in the workspace within the application. However, conversations can be exported to use elsewhere if the user was granted the conversation exporting permissions. See [Exporting a conversation on page 25](#) and [Permissions on page 10](#) for more details.

Conversation history in the panel is also saved across sessions and retained as long as the workspace exists, unless the conversation is deleted. For details on managing conversations, see [Using Conversation Manager](#).

Note: All conversations, including deleted ones, are retained per workspace for defensibility reasons, and can be retrieved by contacting Relativity Support for assistance.

Follow-up memory is limited to the 10 most recent messages in the current user session. As a conversation gets longer, earlier messages may no longer be referenced once the model reaches its context limit.

5 Intent detection

Intent detection is one of the mechanisms that Assist uses to understand the purpose behind each prompt and respond in the most appropriate way. Instead of simply matching keywords, Assist analyzes the intent of a query to determine whether it should:

- Retrieve and summarize information from indexed documents.
- Ask for clarification when the question is unclear or refers to multiple topics.
- Indicate when a request cannot be fulfilled due to missing data or unsupported actions.

This capability is designed to help Assist provide relevant and context-aware answers, which may contribute to a more natural and efficient chat experience.

How intent detection works

When a question is submitted, Assist automatically evaluates the likely intent and determines an appropriate response approach. Depending on the context and available data, one of the following response strategies is used:

Response strategy	Description	Recommended actions
Retrieval-based answering (using RAG)	When a question can be answered using the indexed documents, Assist retrieves the most relevant information and generates a concise, well-supported response.	No action needed.
Refusal due to missing data	If the system cannot locate relevant information within the indexed dataset, Assist will clearly indicate that no matching content is available.	• Verify that the appropriate index is selected. • Confirm that the expected documents are included in the saved search used to build the index. • Rephrase the question to be more specific (for example, include names, dates, or topics). • If information truly does not exist in the indexed content, no further action is required.
Refusal due to unsupported request	When a question asks for something beyond Assist's scope - such as legal opinions, drafting documents, or making recommendations - the system identifies the intent and responds appropriately.	• Rephrase the request as a document-based question (such as, "What documents discuss termination terms?" instead of "Should we terminate the contract?"). • Focus on facts, references, or evidence rather than conclusions or advice.
Disambiguation (clarifying ambiguous queries)	If a question references multiple possible entities or events, Assist will prompt for clarification before continuing.	Provide the requested clarification.
Fallback for Unclear Intent	If the system cannot confidently determine what the question refers to, it will ask a clarifying question. If clarification is not provided, Assist will default to generating a retrieval-based response using the available dataset.	• Add relevant detail such as people, timeframe, or subject matter. • Break broad questions into smaller, more focused inquiries.

Response strategy	Description	Recommended actions
		Confirm whether the goal is to find facts, summarize documents, or surface key themes.

6 Feedback options

Assist provides ways to submit feedback about the responses generated.

6.1 Feedback controls for responses (thumbs up/thumbs down)

The feedback buttons on the response modal allow you to evaluate the quality of the generated response:


- **Thumbs Up**—click the thumbs up icon to indicate the response was accurate, useful, or aligned with expectations. Optionally, enter detailed comments in the text box explaining why you gave it a thumbs up and click **Submit**.
- **Thumbs Down**—click the thumbs down icon to indicate the response was incomplete, incorrect, unclear, or not aligned with expectations. Optionally, enter detailed comments in the text box explaining why you gave it a thumbs down and click **Submit**.

Feedback submitted through these controls is collected in the background for analysis, such as monitoring overall answer quality and identifying recurring issues.

See [Navigating responses](#) for additional information on the response modal.

6.2 In-product feedback survey

You may occasionally be presented with a short feedback dialog within the Assist asking how helpful were Assist's responses for completing your task. Options include: Not helpful, Slightly helpful, Moderately helpful, Helpful, and Very helpful. Your choice selection to this quick question helps us understand how Assist supports real-world workflows and where it delivers the most value.



Assist Feedback

1. How helpful were Assist's answers for completing your task?

☐ 1 - Not helpful

☐ 2 - Slightly helpful

☐ 3 - Moderately helpful

☐ 4 - Helpful

☐ 5 - Very helpful

7 Using Conversation Manager

Conversation Manager provides a central place to organize and navigate all your conversations. From here, you can rename, delete, pin, and export conversations, as well as switch between threads or pick up where you left off in a previous one.

Users must be granted Prompting permission to manage their conversations. See [Permissions](#) for details on the Prompting permission.


Note: ARM is not supported for Assist. Assist permissions, indexes, conversations, and metadata mapping cannot be archived, restored, or moved using ARM.

Suggested resources:

- [Conversation history and sessions on page 20](#) for additional information.
- [Frequently asked questions on page 36](#) for helpful answers to common questions about conversations.

7.1 Navigating the Conversations panel

The Conversations panel slides open next to the Assist chat panel. Refer to the following to learn more about the panel.




- The newest conversation appears at the top of the **Recents** list. Click the drop-down arrow to sort by Newest first, Oldest first, A→Z (for ascending order), or Z→A (for descending order).
- Pinned conversations appear at the top of the Conversations pane in the **Pinned** section.
- Click on a conversation to view it in the Assist chat panel. A vertical highlight bar displays next to a conversation to indicate it is open in the chat panel.
- Hover over the conversation to:
 - View the full conversation name and number of messages in the conversation.
 - View the **More** icon (vertical ellipses) to rename, pin, export, or delete the conversation.
- Click the **X** to close the entire Conversations panel.

7.2 Opening conversation manager

Use the steps below to open conversation manager.

1. To display the Conversations panel, select **Conversations** from the **hamburger** icon in the Assist chat panel.



The Conversations panel slides open next to the Assist chat panel.

2. To close the panel, click the **X** opposite the Conversations panel heading or select **Conversations** again from the **hamburger** icon in the Assist chat panel.

7.3 Viewing a conversation

Use the steps below to open a particular conversation to view its contents.

1. Open the Conversation panel. See [Opening conversation manager on the previous page](#) for details.
2. Locate and click the conversation name.
A vertical highlight bar displays next to the conversation that is open in the Assist chat panel.
3. Optionally, continue asking questions within the conversation or click another conversation name to switch conversations.

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

7.4 Renaming a conversation

By default, when you open Assist initially or begin a new conversation, it automatically names the conversation using the current date (for example, "Conversation MM/DD/YYYY"). After you submit your first prompt or question, the conversation name updates to reflect that content. You can rename any conversation from the Conversations panel or while viewing the conversation.

1. From either the Conversations panel or within an open conversation, click the **More** icon (vertical ellipses) next to the conversation to be renamed.
2. Select **Rename**.
3. Enter new name for the conversation (up to 50 characters) in the **Conversation name** field.



Rename Conversation

Conversation name

Conversation 1/27/2024

Save Cancel

4. Click **Save**.

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

7.5 Deleting a conversation

You can permanently delete an entire conversation from either the Conversations panel or within the active conversation itself. Conversation history is kept as long as the workspace exists, unless the conversation is deleted.

1. From either the Conversations panel or within an open conversation, click the **More** icon (vertical ellipses) next to the conversation to be deleted.
2. Select **Delete**.
3. Click **Delete** on the confirmation dialog to proceed with deleting the conversation.


The conversation permanently disappears from the panel.

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

7.6 Pinning a conversation

For quick access to important conversations, you can pin them to the top of the Conversations list. They will appear above the All conversations section.

1. From either the Conversations panel or within an open conversation, click the **More** icon (vertical ellipses) next to the conversation to be pinned.
2. Select **Pin**. The conversation moves to the Pinned section at the top.



3. To remove a pinned conversation from the Pinned section, click the **More** icon (vertical ellipses) next to the conversation to be unpinned and select **Remove from pins**.

The conversation will return to the All conversations list in date order.

7.7 Exporting a conversation

Users with permission to export conversations can save a conversation as a Microsoft Word (.docx) file while preserving its context, citations, and source links. This is particularly valuable for legal and review teams who need to share findings with external parties (clients, co-counsel, opposing counsel) or internal stakeholders who lack Relativity access. See [Permissions on page 10](#) for more details on setting user permissions.

1. From either the Conversations panel or within an open conversation, click the **More** icon (vertical ellipses) next to the conversation to be exported.
2. Select **Export**.

The exported file is saved to the user's default browser download folder.

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

8 Management

The Index Manager contains the list of indexes created in Assist or Case Strategy. An index is a set of documents that Assist searches to answer prompts. Each index is built from the extracted text of documents available in public saved searches within a workspace. Users can choose the index that best matches their prompt, or switch indexes to search for different topics or document sets in the same workspace.

Note: Indexing and prompting do not incur additional costs.

Users responsible for managing indexes must be granted the Index Management permission to create, rebuild, and delete indexes. Individuals with Prompting permission may access the Index Manager list to select an index for their prompt. See [Permissions](#) for details on the Index Management and Prompting permissions.


Note: ARM is not supported for Assist. Assist permissions, indexes, conversations, and metadata mapping cannot be archived, restored, or moved using ARM.

Suggested resources:

- [Working with indexes on the next page](#) for how to create, edit, delete, export, and maintain indexes.
- [Best practices](#) for recommendations on preparing your data for indexing.
- [Frequently asked questions on page 36](#) for helpful answers to common questions about indexes and index management.

8.1 Indexes created in Case Strategy

A Case Home index can be created within Case Strategy as part of the facts generation process (shown below). It enables Assist to support case strategy development by leveraging the same document set used in Case Strategy analysis. See Case Strategy documentation for more information on this process.



Notes:

- The Case Home index cannot be rebuilt or deleted from Index Manager because it is managed by the Case Strategy application.
 - Metadata mapping is not supported for the Case Home index. It is only supported for indexes created using Assist. See [Metadata Mapping](#) for more information.
-

8.2 Indexes created within Assist

Indexes can be created manually within Assist from public saved searches within the workspace.

- Each index can include up to 300,000 documents.
- A maximum of five (5) built indexes can exist at one time per workspace. This can be from a combination of the Case Home document set (Case Strategy) and public saved searches. See [Working with indexes on the next page](#) for details on adding and deleting indexes.
- The indexes define the datasets from which Assist retrieves and generates answers.

8.3 Throughput expectations

- Indexing speed varies by dataset size, document complexity, and available workspace resources. While typical cases are optimized, larger or complex datasets may take more time, possibly up to 6 hours, to finish with either a "Successfully build" or "Partially build" status.
- Index creation is a background process that progresses automatically once initiated. While exact durations vary, it is recommended to allow sufficient processing time for workspaces containing large volumes of documents or complex extracted text.

9 Working with indexes

Users with the permission to use Index Manager can create, rename, delete, and rebuild indexes. See [Permissions](#) for details on the Index Management permission. A maximum of five (5) indexes can exist at a time per workspace. If the maximum is reached, one of them must be deleted before adding a new one.

Note: ARM is not supported for Assist. Assist permissions, indexes, conversations, and metadata mapping cannot be archived, restored, or moved using ARM.


Suggested resources:

- [Best practices](#) for recommendations on preparing your data for indexing.
- [Frequently asked questions on page 36](#) for helpful answers to common questions about indexes and index management.

9.1 Navigating the Index Manager

Refer to the following descriptions to learn more about the Index Manager modal.

Note: Multiple users can use Assist in the same workspace and use the same indexes, but each user's session and chat history remain private and separate.



- **Search index box**—Use to search for a particular index in the list. Begin typing the name and matching indexes display for selection.
- **Index name**—name of the index entered during index set up. Hover over or click the index name to view index details, such as last indexed date/time and by whom, description, and whether all documents were indexed or errors were encounter.
- **Docs**—total number of documents in the index. When indexing is in progress, the percentage complete displays here.
- **Last indexed**—date the index was last indexed.

- **Create index**—click to create a new index (plus sign icon). See [Creating a new index below](#) for more information.
- **Rebuild index**—click to rebuild (counter-clockwise circular arrow icon) the selected index. See [Rebuilding an index on page 30](#) for more information.
- **Delete index**—click to delete (trash can icon) the selected index. See [Deleting an index on page 31](#) for more information.
- **Set index**—click button after selecting an index to use for querying.

9.2 Indexing rules and limitations

To ensure quality and performance, the indexing process follows specific rules:

- **File size limit:** Documents with more than 5 MB of extracted text are automatically skipped during indexing.
- **Text extraction requirement:** Only documents with extracted text are included.
- **Content used for generation:** Assist uses only the extracted text from indexed documents when retrieving information and generating responses.


9.3 Creating a new index

Use Index Manager to create indexes. A maximum of five (5) indexes can exist at a time per workspace. If the maximum is reached, one of them must be deleted before adding a new one.

Note: Indexing and prompting do not incur additional costs.

To create a new index using Index Manager:

1. Click the **Assist** icon in the side bar to open its chat panel.
2. From the chat panel, navigate to the Index Manager by clicking the index name in the **Ask a question** box.



3. Click the **Create Index (+)** icon in the Index Manager list.

Index Manager ×

Search index...

Index Name	Docs	Last indexed
✓ 4k Docs Index	3997	8/27/2025
✓ 1k docs index	1000	7/17/2025
✓ Case Home	1997	7/7/2025

+

Set index

4. Fill in the following fields:

Create aiR Assist Index ×

Source *

Select a data source

Name *

Description (optional)

Cancel

Create Index

- **Source**—choose a public saved search as the document source for your index. The index will be built from the extracted text of documents in that saved search. Only public saved searches are available for use.
- **Name**—enter a unique, clearly descriptive name for the new index so its contents and purpose are identifiable by any user. Must be at least five characters long.
- **Description (optional)**—optionally add a description of the index, such as its use case or intended function. This information may assist other users in the workspace who might use it.

Note: Assist indexes are shared by all Assist users within the workspace, so keep that in mind as you name and describe each one.

5. Click **Create Index**.

Indexing progress appears below the new index name. See [Working with indexes on page 27](#) for information.

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

You can access other indexes or create new ones while the index is building.

Notes:

- The **Ask a question** box is not available while documents are being indexed.
- Assist indexes are shared by all users within the workspace.


9.4 Rebuilding an index

If your data source (saved search) changes by either adding or removing documents to it, the index must be updated (rebuilt) to reflect those changes for Assist to use the new information in its queries. Additionally, any changes made to metadata mapping, including enabling or disabling metadata, require rebuilding indexes so updates are reflected in the data used for analysis. See [Metadata mapping on page 34](#) for more information.


Note: The Case Home index cannot be rebuilt or deleted from Index Manager, because it is managed by the Case Strategy application.

To rebuild an index from the Index Manager:

1. Click the **Assist** icon in the side bar to open its chat panel.
2. Navigate to the Index Manager from the chat window by clicking the index name in the **Ask a question** box.



3. Select the desired index to rebuild.



4. Click the **Rebuild** curved arrow icon. The index cannot be used during rebuilding, but you can work in other indexes or create a new one.
5. Click **Rebuild** to proceed.

Indexing progress appears below the new index name. See [Monitoring the index build process on page 32](#) for information.

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

9.5 Editing an index name

Use the steps below to edit the name of an existing index. This action will not trigger a rebuild.

To edit an index name from the Index Manager:

1. Click the **Assist** icon in the side bar to open its chat panel.
2. Navigate to the Index Manager from the chat window by clicking the index name in the **Ask a question** box.
3. Select the desired index to edit and click the **Edit** (pencil) icon.

Index Manager

Search index...

Index Name	Docs	Last indexed
4k Docs Index	3997	8/27/2025
1k docs index	1000	7/17/2025
Case Home	1997	7/7/2025

Set index

4. Update the **Name** field as needed. The Data Source field is read-only, displaying the current Data Source name.

Note: The Case Home index cannot be renamed.

Edit aiR Assist Index

1. Data Source ⓘ
BTYS Priv Export

2. Name ⓘ
1k docs index

3. Description (Optional)

Save

Cancel

5. Optionally add or edit the **Description** field.
6. Click **Save**.

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.


9.6 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.


Note: The Case Home index cannot be rebuilt or deleted from Index Manager, because it is managed by the Case Strategy application.

To delete an index from the Index Manager:

1. Click the **Assist** icon in the side bar to open its chat panel.
2. Navigate to the Index Manager from the chat window by clicking the index name in the **Ask a question** box.



3. Select the desired index to delete.



4. Click the **Delete** trash can icon.
5. Click **Delete** to proceed with permanently deleting the selected indexes. After the deletion process begins, it cannot be reversed.

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

9.7 Monitoring the index build process

After initiating the index build process, the Index Manager displays indexing progress percentages in the Docs column. You may work in other indexes or start a new one while indexing is in progress. When the index build finishes, the Index Details display the name of the saved search on which the index was built as well as the number of documents indexed.

The index can complete in two possible states:

- **Successfully built**

A check mark within a circle displays next to the Index Name to show a successful build. The index is successfully built when all documents from the selected saved search are error-free. Index details show the name of the saved search used to create the index, the total document count, and a confirmation indicating all documents were indexed. The index is immediately available in Assist.

Index Name	Docs	Last indexed
✓ Issues	24562	10/04/2025
Index details Build on Privileged Communications Saved Search with 32651 documents. All documents were indexed.		

- **Partially built (Document errors encountered)**

An exclamation mark within a triangle appears beside the Index Name to indicate a partial build, signifying that certain documents were not indexed. Index details show the name of the saved search used to create the index, total document count in the saved search that were indexed, and number of documents that remain unindexed. You may click **View errors** to review the list of errored documents or **Try again** to initiate an index rebuild. See [Error handling and retrying index jobs below](#) and [Working with indexes on page 27](#) for more information.

Index Name	Docs	Last indexed
▼ ⚠ EE Traders	248	2/24/2026
Index details Build on EE Traders Saved Search with 251 documents. 3 documents were not indexed.		
↶ Try again ⓘ View errors		

9.8 Error handling and retrying index jobs

Assist provides error messages and offers retry options for indexes that are partially built. If an index fails, you can click **Try again** to rerun it, or you can click **View errors** to investigate the Indexing Error List. For more information, see [Working with indexes on page 27](#).

1. From the Index Manager, click on the name of the index that was partially built to view details.
2. Do any of the following:
 - Click **Try again** to rebuild the index to see if the errors are corrected.
 - Click **View errors** to review the Indexing Error List.
3. When **View errors** is clicked, the Indexing Error List appears next to the Assist pane and covers any application you have open, such as the Documents List. The name of the selected index appears next to the **Not Indexed Documents** label.
 - To filter the list, click one of the indexing error types in the panel above the list or use the **Indexing Error Type** column to select one from the list.
 - To view the full list again, click **Not Indexed** in the panel.
 - To view a document further in the Viewer, click its **Control Number**. Closing the Viewer returns you to the error list.
 - To close the error list and return to the previously opened application, click the **X** on the error list.

Control Number	Artifact Id	File Size	Indexing Error Type	Relativity Native Type	File Name	Unified Title	System Created By	System Created On
1	BTv5_P_0000000955	1081659	No Extracted Text	Internet Mail Message	BTv5_P_0000000955.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
2	BTv5_P_0000000956	1081660	Other Errors	Internet Mail Message	BTv5_P_0000000956.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
3	BTv5_P_0000000957	1081661	Other Errors	Internet Mail Message	BTv5_P_0000000957.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
4	BTv5_P_0000000958	1081662	No Extracted Text	Internet Mail Message	BTv5_P_0000000958.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
5	BTv5_P_0000000959	1081663	Other Errors	Internet Mail Message	BTv5_P_0000000959.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
6	BTv5_P_0000000960	1081664	File Size > 5 MB	Internet Mail Message	BTv5_P_0000000960.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
7	BTv5_P_0000000961	1081665	Other Errors	Internet Mail Message	BTv5_P_0000000961.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
8	BTv5_P_0000000962	1081666	File Size > 5 MB	Internet Mail Message	BTv5_P_0000000962.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
9	BTv5_P_0000000963	1081667	Other Errors	Internet Mail Message	BTv5_P_0000000963.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
10	BTv5_P_0000000964	1081668	File Size > 5 MB	Internet Mail Message	BTv5_P_0000000964.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
11	BTv5_P_0000000965	1081669	Other Errors	Internet Mail Message	BTv5_P_0000000965.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	
12	BTv5_P_0000000966	1081670	Other Errors	Internet Mail Message	BTv5_P_0000000966.EML	Service Account, Relativity	1/9/2026, 3:17 PM C	

Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

10 Metadata mapping

Incorporating metadata mapping helps improve how Assist retrieves documents by allowing the AI model to use structured metadata as part of the retrieval process, in addition to semantic relevance. This leads to more accurate, context appropriate results when answering user questions.

- Without the metadata mapping and usage, Assist often receives questions with constraints not clearly stated in the document, such as sender or recipient email information and date ranges.
- With metadata mapping enabled, Assist can detect these elements in a natural language question and translate them into metadata filters at retrieval time. This ensures that documents considered for the answer align with the user's intent, not just with similar wording.

Assist uses a hybrid retrieval approach when metadata mapping is configured:

1. Metadata filters are applied first, based on detected entities such as emails and dates.
2. Semantic similarity search is then performed on the filtered set of documents.

By restricting semantic search to documents that already meet the relevant metadata criteria, Assist helps to reduce noise and improves ranking quality. This approach is more precise than relying on semantic search alone. For more information on Assist's usage of a Retrieval-Augmented Generation (RAG) architecture, see [Understanding Assist responses on page 7](#).

Without metadata filtering, semantically related but contextually incorrect documents (for example, the wrong time period or sender) may be included in retrieval. This can negatively impact answer quality. Metadata mapping helps ensure that Assist reasons over documents that are valid for the question being asked, reducing the likelihood of incorrect or over generalized responses.

Assist supports five pre-configured mapping fields: Primary Date/Time, Email From, Email To, Email CC, and Email BCC. New ones cannot be added.

Note: ARM is not supported for Assist. Assist permissions, indexes, conversations, and metadata mapping cannot be archived, restored, or moved using ARM.

10.1 Considerations


- Each workspace can have only one metadata mapping configuration and it applies to all indexes listed in Index Manager.
- Metadata mapping is only supported for indexes created using Assist. It is not supported for the Case Home index, which is managed by Case Strategy.
- Metadata mappings included in the template are not copied to new workspaces created from that template.

10.2 Configuring metadata mapping

Workspace administrators can configure how standard metadata fields are mapped to their specific workspace-related fields, with field names and IDs uniquely assigned to each workspace.


Note: Any changes to metadata mapping, including enabling or disabling metadata, require rebuilding indexes so updates are reflected in the data used for analysis. See [Rebuilding an index on page 30](#) for more information.

1. Click the **Assist** icon in the side bar to open its chat panel.
2. Select **Metadata Mapping** from the hamburger icon in the Assist chat panel.



The Metadata Mapping panel slides open next to the Assist chat panel. To close the panel, click the **X** opposite the panel heading or click **Metadata Mapping** again from the hamburger icon.

3. Toggle **Enable Metadata Mapping** to set up and use metadata in the workspace. By default, the toggle is disabled.



The Fields Mapped list is pre-configured.

4. For each relevant metadata field listed that you want to use, select the appropriate workspace-level field from the corresponding drop-down. The fields listed in the drop-down will vary between workspaces. The first three metadata fields must be set up: **Primary Date/Time**, **Email to**, and **Email from**. Maximum of five fields total can be configured. Labels appear next to each metadata field indicate the mapped or not mapped status.

Field	Field type	Required field	Description
Primary Date/Time	Date/Time	Required	Used to understand when document was created or sent in queries that require date/time filters.
Email To	Long text	Required	Used to identify the email addresses to which the email was sent in queries pertaining to emails or specific participants.
Email From	Fixed-length text	Required	Used to identify the email address that sent the email in queries that ask about emails or specific participants.
Email CC	Long text	Optional	Used to identify the email addresses included as carbon copy (CC) recipients in queries pertaining to emails or specific participants.
Email BCC	Long text	Optional	Used to identify the email addresses included as blind carbon copied (BCC) in queries that ask about emails or specific participants.

5. Click **Save & update indexes**.

6. Click **Confirm** to proceed with mapping metadata fields. This action automatically updates all indexes in the workspace.

The re-indexing process may take some time to complete. To view its progress, access the Index Manager. For more information, see [Monitoring the index build process on page 32](#).


Audit logs capture this action. See [Auditing user activity on page 8](#) for more information.

10.3 Using metadata for prompting

The **Use Metadata** toggle displays in the prompting Ask a question box if Metadata Mapping is enabled. Each user can choose whether or not to leverage metadata for their inquiry by toggling **Use Metadata** on or off. See [Prompting](#) for details.

Notes:

- Users with Prompting permissions enabled can use this toggle. See [Permissions on page 10](#) for more information.
- The Use Metadata toggle is only supported for indexes created using Assist. Currently, Case Strategy Case Home index does not support metadata mapping so the toggle is disabled.



11 Frequently asked questions

Listed below are questions you might have using the Assist application.

11.1 Prompt FAQs

What types of prompts can I ask in Assist?

Assist is designed to respond to prompts related to the documents in the current workspace and not general knowledge prompts. Write explicit and concise prompts using natural language, such as "Who are the key people in

this case?" or "Which documents mention [topic]?" For more guidance on writing prompts, see [Best practices on page 12](#).

Is there a limit to how many prompts I can ask in one session?

There is no limit to the number of prompts allowed, but long or highly complex queries may take more time to process. For best results, keep your prompts concise and focused. For more guidance on writing prompts, see [Best practices on page 12](#).

Does Assist ever reject my prompts?

Yes. Assist uses advanced intention capabilities to reject prompts that cannot be answered or are outside its scope. It also requests clarification for vague or overly broad queries. See [Intent detection on page 20](#) for more information.

Are there recommended prompt patterns or techniques to improve accuracy for contract or relationship extraction tasks?

Assist uses a Retrieval-Augmented Generation (RAG) approach to respond to prompts. The system first retrieves the most relevant documents for the query and then uses a large language model (LLM) to generate a summarized response. The following best practices can help improve accuracy when working with contract, entity, or relationship extraction tasks. For the full list, see [Best practices on page 12](#):

- Be explicit, not implicit
Retrieval systems work best with clearly defined prompts. Vague prompts, such as "Tell me about the case," are less effective than specific ones, like "Describe contract terms related to indemnification."
- Write concise prompts
Multi-part or compound prompts can fragment retrieval results. Breaking complex prompts into smaller, focused ones improves precision and clarity.
- Leverage keywords and synonyms
Retrieval engines benefit from varied phrasing. Include alternative terms or related entities, such as "bribe," "gift," or "incentive," to capture a broader range of relevant content.
- Lack of comprehensiveness
Retrieval systems aim to respond to the prompt directly, rather than provide a full list of all related documents or entities.

Can Assist understand multiple languages?

The large language model (LLM) used by Assist has been evaluated for 83 languages. Although Assist has primarily been tested on English-language documents, it is designed to support non-English datasets. For more information on this issue, see [Language support on page 7](#).

Can I share individual prompts asked in Assist with others?

No. You cannot share individual prompts with others within Assist. However, full conversations can be exported to an MS Word (.docx) file without losing the context, citations, or source links and shared with others that way. See [Exporting a conversation on page 25](#) for details.

How can I limit users to only asking prompts, without permission to create, rebuild, or delete indexes?

You must set up their permissions to only allow access to Prompting but not to Index Management. Similarly, if you have users who are only responsible for managing indexes (create, rebuild, delete), you could set them up with only Index Management permission, but not Prompting. See [Permissions on page 10](#) for details.

Am I able to bookmark important prompt and response pairs within a conversation thread so that I can quickly and easily find them?

Yes. You click the star icon below a response. See [Starring \(bookmarking\) a prompt and response on page 20](#) for details.

Is it possible to pin an entire conversation so that I can quickly access it?

Yes. Important conversations threads can be pinned to the top of the Conversations list. They will appear under the Pinned section. See [Pinning a conversation on page 25](#) for details.

11.2 Response FAQs

Can Assist be configured to produce more consistent or deterministic responses?

No. Assist does not currently include configuration options to control response variability or “thinking time.” The system used by Assist is designed to reduce variability, but some level of randomness is inherent to all large language models (LLM). As a result, identical prompts may generate slightly different responses, especially when phrased implicitly or influenced by prior conversation context. The more specific and well-defined the prompt, the more consistent the output tends to be. Assist focuses on accuracy and relevance rather than adjustable response styles. For more information, see [Understanding Assist responses on page 7](#).

Does Assist filter out or de-prioritize repeated or boilerplate content (such as disclaimers or email signatures)?
Not at this time. Assist follows a Retrieval-Augmented Generation (RAG) approach, in which only the most relevant document chunks are retrieved and used to generate a response. Repeated or boilerplate text (such as, standard disclaimers, email footers, or signatures) is not explicitly filtered out. However, this type of content is typically not selected unless it is determined to be highly relevant to the user’s query.

How is the accuracy of responses ensured?

Assist relies on indexed data to generate responses leveraging retrieval-augmented generation (RAG). While it aims to offer correct responses, always verify the offered information and citations for accuracy. Using the feedback buttons in the response box can help improve the tool over time. See [Feedback options on page 22](#) for details on using the feedback buttons.

How can I leave feedback (like or dislike) within the Assist application?

To leave feedback in the Assist app, use the thumbs up or down icons below each response to rate it and add comments. See [Feedback options on page 22](#) for more information. You may also occasionally be present with an Assist Feedback dialog asking how helpful were Assist's responses for completing your task. See [In-product feedback survey on page 22](#) for details.

11.3 Conversation FAQs

Is my conversation history private to me or can others see it?

Each user's session and conversation history remains private and no one else can see it within Assist. Additionally, they cannot be shared with or accessed by others in the workspace. However, you can export the contents of a conversation to an MS Word (.docx) file and share them that way. See [Exporting a conversation on page 25](#) for details.

Can I delete any of my conversation history?

Yes. You can delete an entire conversation and its history using Conversation Manager. See [Using Conversation Manager](#) for details.

Does my conversation history ever clear?

Your conversation history is saved across your user sessions and retained as long as the workspace exists, unless the conversation is deleted. See [Using Conversation Manager](#) for details on deleting conversations. Also refer to [Conversation history and sessions on page 20](#) for more information surrounding conversation history.

Can I share conversations with other users in Assist?

No. You cannot share conversations with others within Assist. However, you can export the contents of a conversation to an MS Word (.docx) file and share them that way. See [Exporting a conversation on page 25](#) for details.

Can I change the name of the conversation?

Yes. You can rename an entire conversation using Conversation Manager. See [Using Conversation Manager](#) for details.

Can I go back to a past conversation to enter more prompts?

Yes. You can switch between conversations anytime, as well as return to a previous conversation to continue querying using Conversation Manager. See [Using Conversation Manager](#) for details.

Can multiple users access Assist simultaneously?

Yes. Multiple users can use Assist within the same workspace concurrently. However, each user's session and chat history remain private and separate.

Does Assist maintain conversational continuity?

Yes. Assist maintains conversational continuity within a session, allowing it to reference the context of previous queries when formulating responses. This means that follow-up prompts can build on information discussed earlier in the same conversation.

However, the continuity mechanism in Assist differs from general-purpose chat systems, such as ChatGPT. Each interaction still begins with a new document retrieval step, where Assist searches for the most relevant materials based on the current query and prior context. The retrieved content then informs the contextualized prompt, which the language model uses to generate the final response.

Is it possible to export the conversation so I can share it elsewhere?

Yes. Conversations can be exported to an MS Word (.docx) file without losing the context, citations, or source links.

See [Exporting a conversation on page 25](#) for details.

11.4 Index FAQs

Where are Assist indexes stored, and how can indexing errors be reviewed?

Assist indexes do not currently have a dedicated view within Relativity aiR. Indexes are managed directly through the Assist interface. When switching between indexes, the system displays a status card within the chat showing the index state.

If an index was built with errors, the status card includes a “View Errors” option that lists documents that encountered issues during indexing. Documents larger than 5 MB of extracted text are automatically skipped during the indexing process and are not currently displayed as errored items.

See [Working with indexes on page 27](#) for more information.

How can I limit users to only manage indexes (create, rebuild, or delete), without permission to enter prompts?

You must set up their permissions to only allow access to Index Management. See [Permissions on page 10](#) for details.

11.5 Document FAQs

Is Assist able to read Excel files, including multiple tabs, formulas, and embedded objects? Are there any formatting or processing steps we should apply first?

Assist generates its responses based solely on the extracted text from processed documents. This means that for Excel files, the content available to Assist depends on how text is extracted during processing.

Depending on your processing profile settings and the engine used to extract text, certain elements (such as data from multiple tabs, formulas, or embedded objects) may or may not be included in the extracted text. To ensure the most complete and accurate results, review and optimize your processing configuration before indexing Excel files.

Is Assist able to reliably read documents that contain only images and no native text files?

Assist can only use the extracted text available in Relativity aiR. For image-only documents (such as scanned PDFs or pictures without embedded text), text extraction must be performed in advance, either through an external process or using an OCR (Optical Character Recognition) workflow within Relativity aiR.

Without extracted text, Assist will not be able to index or generate responses from image-based documents.

Are any additional steps required for large documents (under 5 MB)?

No additional steps are required for large documents that are within the 5 MB of extracted text size limit. During indexing, Assist automatically divides large documents into smaller text chunks to optimize retrieval and improve response relevance. These chunks may be retrieved individually when Assist processes a query, ensuring that large documents are handled efficiently without additional configuration.

Are documents over 5 MB included in the error report?

No. Documents exceeding the 5 MB of extracted text size limit are automatically skipped during indexing and are not currently flagged or listed in the error report. As a result, these files will not appear among errored documents in the Assist interface.

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.

©2026. Relativity ODA LLC. All rights reserved. Relativity® is a registered trademark of Relativity ODA LLC.

12 Index

A

Asking a question 15

Asking a question in Assist 15

Assist Feedback 22

C

Chat history 19

F

feedback 22

feedback survey 22

N

New conversation 19

S

Start a new conversation 19