RelativityOne

aiR Assist Guide

May 8, 2026

For the most recent version of this document, visit our **documentation website**.

Table of Contents

1 aiR Assist (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

aiR Assist is a conversational search tool integrated within RelativityOne, designed to empower legal teams to interact with their data using natural language. By leveraging advanced AI, aiR 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, aiR Assist identifies the documents deemed most relevant and employs a large language model (LLM) to generate answers, complete with citations from as many as 25 source documents.

Note: aiR Assist is not available in repository workspaces.

1.1 How aiR Assist works

aiR Assist operates using a Retrieval-Augmented Generation (RAG) process to deliver grounded, evidence-based responses. This approach combines document retrieval with large language model generation to help support

accuracy, transparency, and contextual relevance.

1. **Indexing the documents (indexing step)**
   The user identifies documents to query and creates an index.
2. **Asking a question (question step)**
   The user asks a question.
3. **Finding relevant documents (retrieval step)**
   Each question is matched against the text indexed from the identified documents. aiR Assist performs a similarity search to identify the most relevant content. The documents are divided into smaller passages, and the system selects results that are estimated to correspond most closely to the question.
4. **Generating the answer (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, including up to 25 citations and references to the original sources.

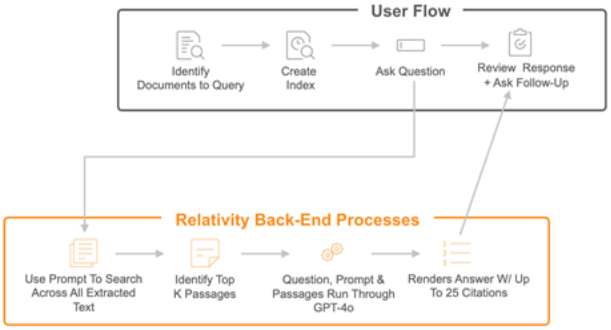

1.1.1 LLM model in use

aiR Assist currently uses the Azure OpenAI LLM, aiming to provide contextually grounded responses based on retrieved source material.

1.1.2 Important Limits

- Each index can contain up to 300,000 documents.
- A maximum of five (5) built indexes can be created per workspace from a combination of the Case Home document set and public saved searches.
- Individual documents must be 5 MB or smaller; larger files are excluded during indexing.
- Only documents with extracted text are indexed. The text must be stored in Data Grid (not SQL). Files that do not contain extracted text are excluded automatically from the index.

1.1.3 Understanding aiR Assist responses

aiR Assist is designed to identify and summarize potentially relevant information from large document sets through natural language interaction. The system operates on a Retrieval-Augmented Generation (RAG) architecture, which retrieves and analyzes the documents most likely to be relevant and generates a response based on retrieved content and supported by citations.

aiR Assist is designed to return contextually relevant and evidence-based information rather than performing exhaustive or “find everything” searches. It does not review every document individually, and some occurrences of keywords or topics 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.

1.2 Language support

aiR Assist currently supports English-language content only. The system has been designed and tested exclusively on English-language datasets to ensure accuracy, reliability, and consistent performance.

At this time, non-English languages are not supported, and aiR Assist has not been formally evaluated or validated for use with multilingual or non-English text. While it may operate with non-English datasets, results can vary in accuracy and completeness, and verification of cited sources is strongly recommended when working with such content.

Future updates may expand language capabilities based on performance testing and model availability.

1.3 Supported use cases

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

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

1.4 Release notes

This section includes the release information and the current functionality of the aiR Assist (Advanced Access) application.

Release: May 4, 2026

- **Feature Permissions**—System Administrators now have the ability to restrict access to aiR Assist per group within a workspace using two Feature Permissions: Index Management and Prompting.

Release: April 17, 2026

- **Metadata Mapping**—ability to use metadata fields (Primary Date, Email From/To/BCC/CC) during indexing to improve how aiR Assist finds and prioritizes documents.
- **Reasoning summary**—ability to see live updates in the interface as aiR Assist works on your question, including the searches being run and a summary of its reasoning.

Release: March 31, 2026

- **Conversation Manager**—ability to create, name, and delete conversations, as well as switch between various threads and resume prior discussions.

Release: February 25, 2026

- **Indexing Progress**—new feature displaying percentage progress of indexing build or rebuild in the Docs column in Index Manager.
- **Indexing Data Source**—added Saved Search name to Index Description in Index Manager.

Release: February 3, 2026

- **Indexing Error List**—new way of displaying additional information regarding not indexed documents including the reason why a document failed to index.

Release: October 20, 2025

- **New aiR Assist button**—aiR Assist is available from the left sidebar instead of under the "Ask AI" button.
- **New chat user interface**—improved chat experience with the aiR Assist chat panel now opening on the left, while the Documents list is always in view on the right.
- **Relocation of start a new chat**—the start a new chat conversation option moved to the Answer a question box.
- **New Index Manager**—ability to filter Indexes and Data Sources (when creating an Index).
- **Answer sharing**—ability to copy answer to clipboard.
- **Give Feedback**—ability to rate an answer with Thumbs Up or Thumbs Down icons.

Release: July 24, 2025

- **Index rebuild**—existing indexes can now be quickly updated to match changes in the saved search data source.
- **Intent clarification**—improves aiR Assist conversations by asking follow-up questions when your request is unclear, helping guide you to the right outcome and navigate the limitations.

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 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" box**—when the 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.

2 Installation (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

aiR Assist does not need to be installed. It is available in all workspaces with aiR for Case Strategy enabled, provided that a Fact Extraction or Transcript Summary job has been completed.

The proper permissions must be configured to use aiR Assist. See Permissions for more information.

Notes:

- aiR Assist requires active use of aiR for Case Strategy.
- aiR Assist is not available in repository workspaces.

3 Permissions (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

System administrators can control access to aiR Assist per group within a workspace using Features Permissions for Index Management, Prompting, and Access Management. 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: aiR Assist permissions set within a workspace will not be retained when restoring the workspace from an archive using ARM.

Below are some typical scenarios, along with the appropriate permission toggle settings:

- Group A can manage indexes, ask questions/prompts, and manage their conversations: Index Management=ON, Prompting=ON, and Access Management=OFF.
- Group B is only responsible for creating, rebuilding, and deleting indexes via Index Manager: Index Management=ON, Prompting=OFF, and Access Management=OFF.
- Group C is only responsible for asking questions (prompting) as well as managing their conversations: Index Management=OFF, Prompting=ON, and Access Management=OFF.
- Group D is only responsible for setting up permissions for other user groups and cannot use any aiR Assist functionality: Index Management=OFF, Prompting=OFF, and Access Management=ON.

Use the steps below to set feature permissions per group within a workspace:

1. Navigate to **Workspace Details** within the **Admin** tab.
2. Click **Manage Workspace Permissions** in the **Relativity Utilities** console.
3. Click the **Features** tab.

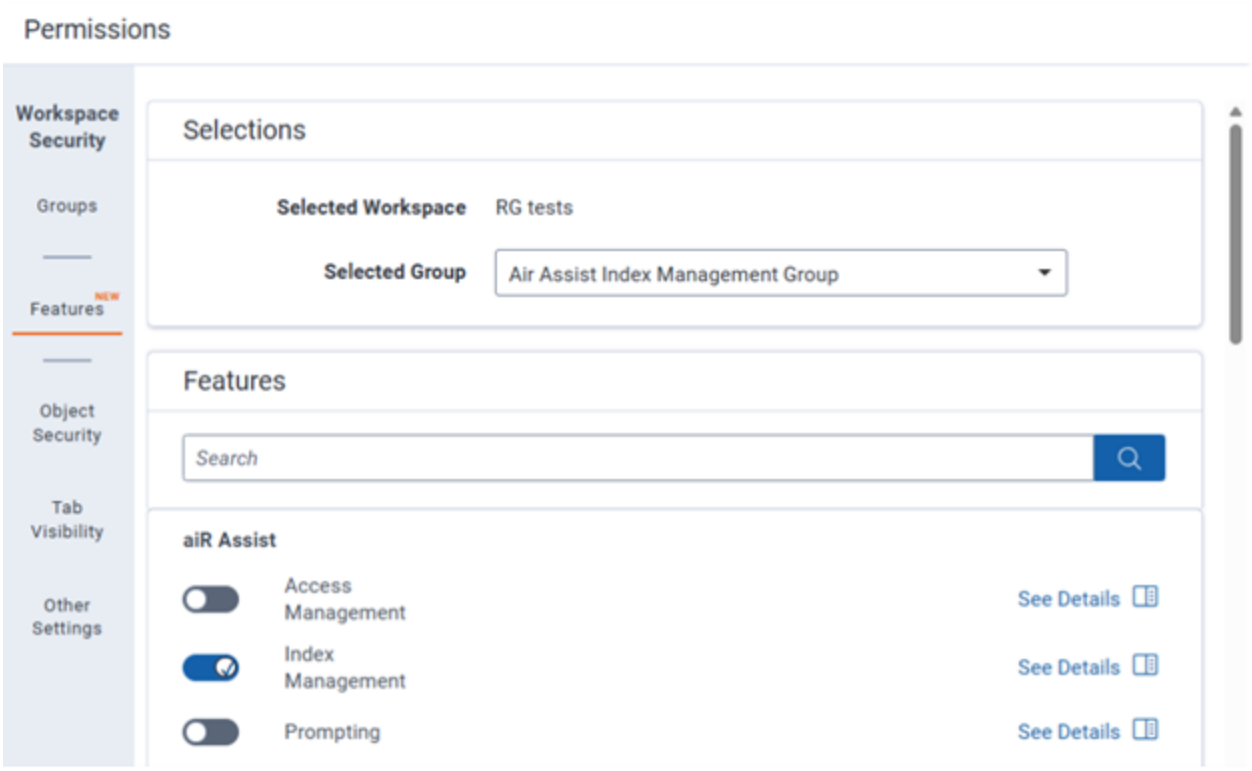

4. Set the aiR Assist permissions as needed for the group:

   Note: When *both* Index Management and Prompting permissions are disabled, aiR Assist will not appear within the group's workspace. If at least one of these permissions is enabled, aiR Assist will appear.

   - **Access Management**—Toggle this setting ON to allow group members the ability to assign and manage the Index Management and Prompting permissions for other groups in the workspace. If no additional aiR Assist permissions are enabled for the group, aiR Assist will not display to them within RelativityOne.

     Note: Given the advanced nature of this setting and its impact on group permission assignment, it is recommended that System Administrators enable it for select groups only.

   - **Index Management**—Toggle this setting ON to allow group members to use the Index Manager to create, rebuild, and delete indexes and use metadata mapping in the workspace. If this is the only permission toggled on for the group, users will not be able to ask questions/prompts or start and manage conversations.
   - **Prompting**—Toggle this setting ON to allow group members to ask questions/prompts about documents in the workspace within created indexes. Users will also be able to use Conversation Manager to create, edit, and delete their conversations.
5. Click **Review and Save** to review the unsaved changes. For more information, see Setting workspace permissions.
6. Click **Save**.

4 Best practices (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

Refer to the best practices below to effectively use aiR Assist.

4.1 Prompting

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

1. **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.
2. **Leverage keywords and synonyms**
   Retrieval engines like specific terms. Include likely variations of keywords (such as, bribe, gift, incentive) or entity synonyms (such as, bt or bt.us for Big Thorium). Retrieval benefits from alternative phrasing.
3. **Stay within the saved search context**
   aiR 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 aiR Assist can draw from within the workspace. Questions should therefore relate to the content of those saved searches rather than general or external topics.
4. **Expect some variation in repeated queries**
   Submitting the same or similar questions multiple times may produce slightly different answers, as aiR Assist regenerates responses dynamically. However, the core content and conclusions are generally expected to remain similar, though variation may occur.
5. **Review citations and supporting references**
   Each aiR 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.
6. **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.
7. **Avoid overly broad or “find everything” queries**
   aiR 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 aiR Assist or, depending on the use case, consider using our other aiR Suite products.
8. **Avoid Using aiR Assist for Calculations**
   Don’t rely on aiR Assist to add up invoices, total damages, or perform complex calculations. The right data may be scattered across multiple files, and LLMs aren’t perfectly reliable at maths.
9. **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 aiR Assist retrieves the most relevant information and clearly understands your objective.

10. **Don't Rely on Metadata for filtering or context**
    Currently, aiR Assist works on extracted text, not metadata. If you need to filter by metadata (such as, only emails between specific dates), results may be incomplete.

Here are examples of some supported and unsupported prompts:

Supported prompt examples

- What evidence is there of [issue]?
- Did [actor] discuss [topic]?
- What topics did [actor] discuss with [actor]?
- Who is [actor]?
- What is [actor]'s role in the company?

Unsupported prompt examples

- Find me **all** examples of [issue].
- Find me **all** conversations that took place between [actor] and [actor].
- Find me emails between March 1-March 31, 2012.
- Find me emails sent after-hours.
- Write me a prompt for aiR for Review related to this matter.
- What is [document name] about?
- What are the discrepancies between [doc 1] and [doc 2].

4.2 Indexing

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

- Verify that documents contain extracted text to ensure they can be properly indexed and used by aiR Assist.
- Verify that public saved searches are well-defined and include only relevant materials.
- Exclude excessively large documents that may be skipped during indexing, in addition to duplicate and known irrelevant documents.
- Ensure that the Extracted Text field is Data Grid File System enabled, as aiR Assist does not support environments where the text is stored in SQL.
- ARM (Archive, Restore, and Move) is not supported for aiR Assist indexes.

5 Prompting (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

You can prompt aiR Assist about indexed documents using conversational language. Asking follow-up prompts builds on the previous one. aiR Assist follows the context of the conversation and people or events discussed. Refer to Best practices (Advanced Access) on the previous page for tips on writing prompts.

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

Note: Prompts and responses are private to each user and persist across the user's sessions. See Using Conversation Manager more information on managing conversations.

5.1 Navigating the prompt composer

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

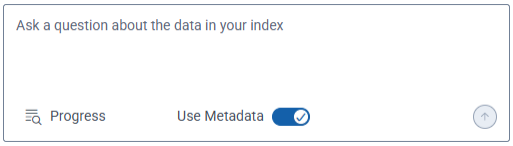

- **Ask a question**—type your prompt into this box. Prompts must be at least three characters long to submit.
- **Index used**—name of the index used for the query appears in the lower left corner of the box. Hover over the name to view the saved search source used, number of documents within the index, date/time last indexed and by whom, and description entered when the index was created.
- **Use Metadata** toggle—if metadata 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 (Advanced Access) on page 26 for more information.
- **Send message**—after typing your question, click the circled up arrow to submit it for a response.

5.2 Navigating the responses

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

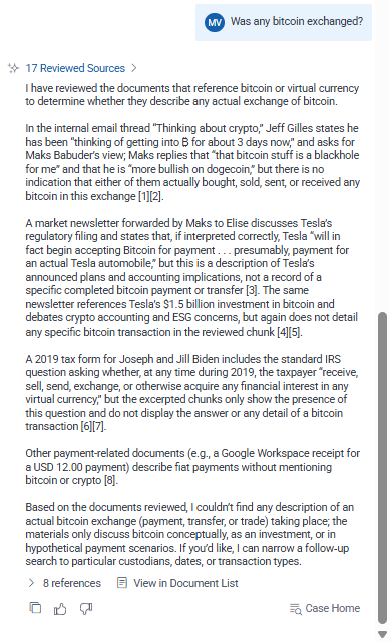

- **Question asked**—the question asked, along with the user's initials, displays above the response.
- **# Reviewed Sources**—click to view the summary, which outlines the steps the AI agent took to answer the question.
    - It includes the name of the index (data source) used, what searches it attempted, potentially relevant documents identified, the documents identified as highly relevant, and so forth. Where available, click the link to view specific documents of that step.
    - The "Attempted following searches" step displays the metadata filters the model applied. During review of the summary, if it is determined that the model selected inappropriate metadata filters, we recommend adjusting the prompt accordingly or disabling metadata usage for the query (as in, disable the **Use Metadata** toggle in the Ask a question box).

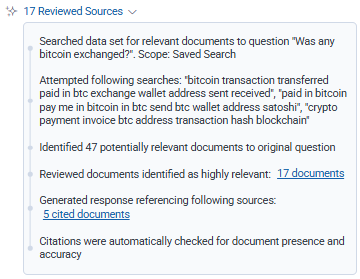

- **Natural language response**—the answer to the question based on the retrieved documents, with inline citations to specific document references.
- **# References**—click the arrow next to the number of references to show or hide the full list of referenced documents generated for the response.
    - aiR Assist can return up to 25 references per question.
    - Below the **Document Control Number** link are the first few pieces of information associated with the reference. Click the down arrow to view more information.
    - Click the Document Control Number link to open the document in the Viewer.

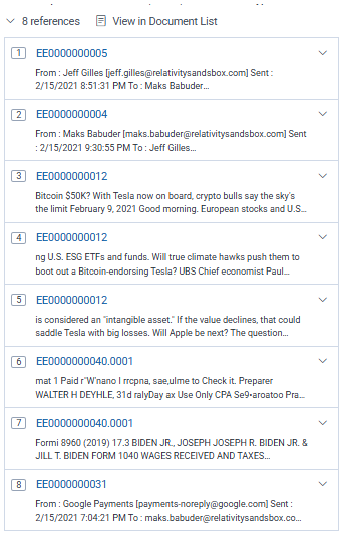

- **View in Document List**—click to filter the Document List to only show documents that include a reference. Be aware that choosing a specific folder in the Folder Browser might change which references appear in the Document List. To ensure all references are visible, select the root folder level.

- **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.
- **Action buttons**—click the action buttons to do the following:
  - **Copy icon**—click to copy the entire response displayed to the clipboard, including citations. See Copying a response on the next page 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 18 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 18 for more information.

See Understanding aiR Assist responses on page 4 for more details on how the system operates on a Retrieval-Augmented Generation (RAG) architecture to retrieve and analyze documents for responses.

5.3 Submitting prompts

To submit a prompt:

1. Click the **aiR Assist** icon in the side bar to open the aiR Assist chat panel.

2. 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**.
3. Type your prompt in the **Ask a question** box using natural language, as if chatting with a person.

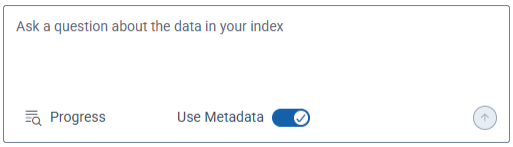

4. If your workspace has metadata enabled, the **Use Metadata** toggle is enable by default. Choose whether or not to leverage metadata for the prompt. See Metadata mapping (Advanced Access) on page 26 for more information.

   Note: The Use Metadata toggle is only supported for indexes created using aiR Assist. Currently, aiR for Case Strategy indexes do not support metadata mapping so the toggle is disabled.

5. Click the **Send Message** up arrow icon to submit the question.
   - aiR Assist displays a response that includes a list of references and documents with control numbers. The application can return up to 25 references in response to a single question.
   - We recommend reviewing the accuracy of the answer and its citations. For additional information on responses, see Navigating the responses on page 11 and Understanding aiR Assist responses on page 4.
6. 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.

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
- persists across each user's Relativity session

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

5.4 Starting a new prompt conversation

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

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

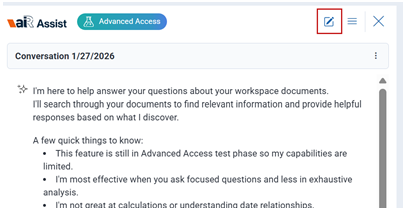

For more information on managing individual conversations, see Using Conversation Manager.

5.5 Copying a response

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

5.6 Conversation history and sessions

aiR Assist keeps a continuous conversation history for each user, ensuring consistency across questions in a session. Each user's conversation data is kept private and is not available to others within the workspace. The conversation history is also saved between user sessions. For more information on managing conversations, see Using Conversation Manager

It is important to note that this conversational memory for follow-up questions is limited by the context window size of 10 previous messages within the current user session. As the conversation grows longer, earlier parts of the exchange may no longer be referenced in responses once the model's context capacity is reached.

6 Using Conversation Manager (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

Conversation Manager lets you access and manage all your conversations. You can see a list of your conversations, rename them, or delete any unneeded threads. It also enables you to switch between different conversations and return to previous ones to continue your inquiries. See Conversation history and sessions above for additional information.

Users must be granted Prompting permission to manage their conversations. See Permissions for details on the Prompting permission.

6.1 Navigating the Conversations panel

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

- The newest conversation appears at the top of the list.
- The vertical highlight bar next to a conversation indicates that conversation is open in the aiR Assist chat panel.
- The message count displayed beneath the name of the conversation indicates the number of questions asked in that thread.
- Conversations are categorized into these categories: Today, Yesterday, This Week, Last Week, and Previous Weeks. The category, “This Week,” excludes those conversations already listed under “Today” or “Yesterday.” The week starts on Monday and ends on Sunday. The category a conversation falls into is based on the date of the last response.

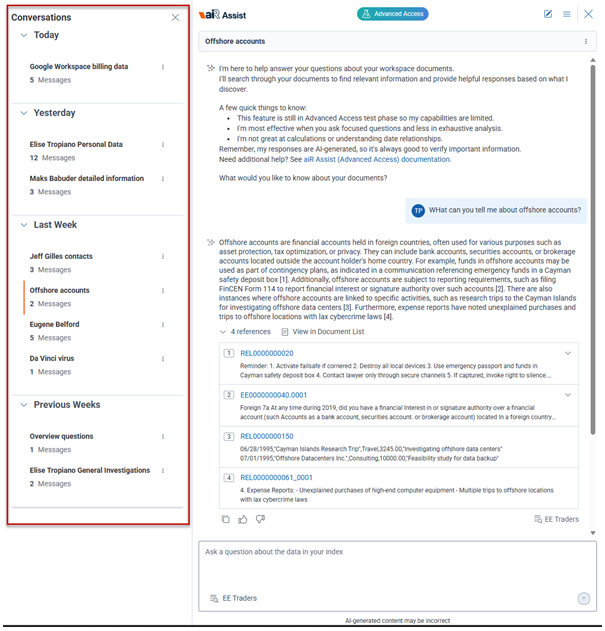

6.2 Opening conversation manager

Use the steps below to open conversation manager.

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

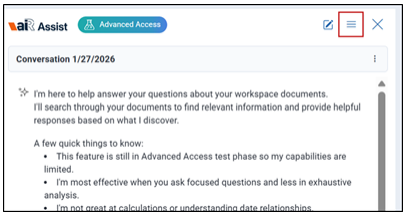

2. The Conversations panel slides open next to the aiR Assist chat panel. To close the panel, click the **X** opposite the panel heading or click **Conversations** from the **hamburger** icon again.

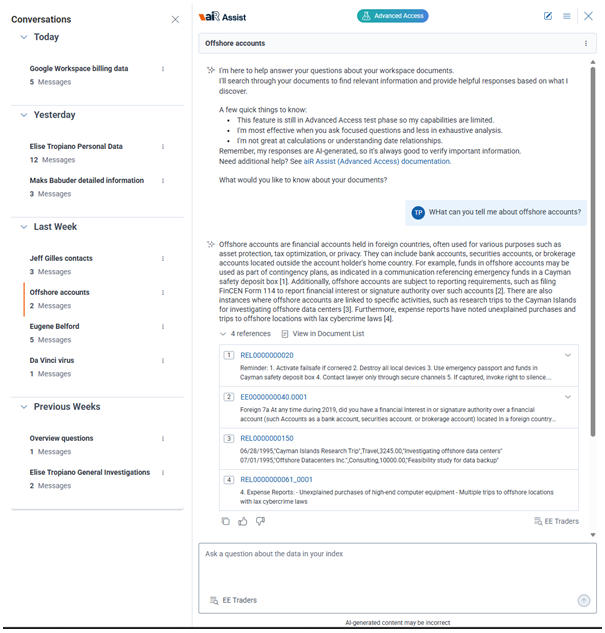

3. To open a particular conversation from the panel, click the **bold conversation name**. The number of messages listed below the name is not clickable. A vertical highlight bar displays next to the conversation that is open in the aiR Assist chat panel.
4. Continue asking questions within the conversation or click another conversation name to switch conversations.

6.3 Renaming a conversation

By default, when you begin a conversation, it is named after the current date (for example, Conversation MM/DD/YYYY). You can change the name of any current or past conversation from either the Conversations panel or within the active conversation itself.

1. From either the Conversations panel or within an active conversation, click the **vertical ellipses** icon next to the conversation to be renamed and select **Rename**.

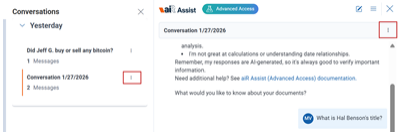

2. Enter new name for the conversation (up to 50 characters) in the **Conversation name** field.

3. Click **Save**.

6.4 Deleting a conversation

You can permanently delete an entire conversation from either the Conversations panel or within the active conversation itself.

1. From either the Conversations panel or within an active conversation, click the **vertical ellipses** icon next to the conversation to be deleted and select **Delete**.

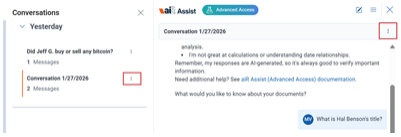

2. Click **Delete** on the confirmation dialog to proceed with deleting the conversation.

The conversation permanently disappears from the panel.

7 Intent detection (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

Intent detection is one of the mechanisms that aiR Assist uses to understand the purpose behind each prompt and respond in the most appropriate way. Instead of simply matching keywords, aiR 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 aiR 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, aiR 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, aiR 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 data-set, aiR 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.

Response strategy	Description	Recommended actions
Refusal due to unsupported request	When a question asks for something beyond aiR 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, aiR 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, aiR 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. • Confirm whether the goal is to find facts, summarize documents, or surface key themes.

8 Feedback options (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

aiR Assist provides ways to submit feedback about the responses generated.

8.1 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 and used to:

- Monitor overall answer quality
- Identify recurring issues

- Improve future model behavior and tuning

See Navigating responses for additional information on the response modal.

8.2 In-product feedback survey

You may occasionally be presented with a short feedback dialog within the aiR Assist asking how helpful were aiR 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 aiR Assist supports real-world workflows and where it delivers the most value.

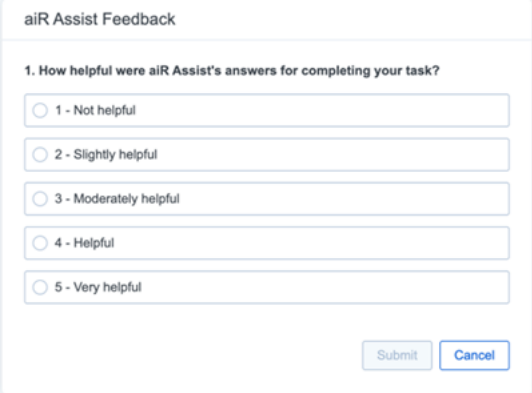

9 Index Management (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

The Index Manager provides centralized access to indexes created in air Assist or aiR for Case Strategy. An index within aiR Assist consists of a collection of documents that the system searches to generate responses to prompts. Each index is constructed from the extracted text of documents available in public saved searches within a workspace. Users may select the appropriate index for their prompts to ensure relevant and informed responses. Additionally, they can switch between indexes when submitting prompts, facilitating exploration of varied subjects or collections within the current workspace.

Designated users who are responsible for index management have the authority to create, rebuild, and delete indexes as needed must be granted the Index Management permission. 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.

Refer to Best practices for recommendations on preparing your data for indexing.

9.1 Indexes created in aiR for Case Strategy

When aiR Assist is used within aiR for Case Strategy, an index can be created either as part of the facts generation process (shown below) in aiR for Case Strategy or independently within aiR Assist. Both methods produce the same result, enabling aiR Assist to support case strategy development by leveraging the same document set used in Case Strategy analysis. See aiR for Case Strategy documentation for more information.

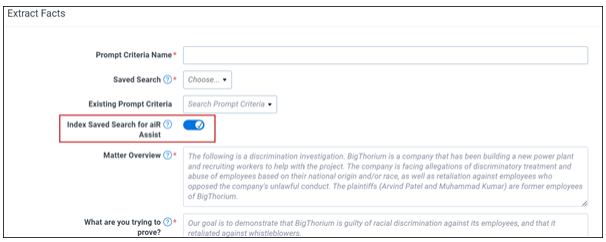

Notes:

- aiR for Case Strategy indexes cannot be rebuilt or deleted from Index Manager because they are managed by the aiR for Case Strategy application.
- Currently, aiR for Case Strategy indexes do not support metadata mapping. It is only supported for indexes created using aiR Assist. See Metadata Mapping for more information.

9.2 Indexes created within aiR Assist

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

- Each index can include up to 300,000 documents.
- A maximum of five indexes can be created per workspace.
- The indexes define the datasets from which aiR Assist retrieves and generates answers.

9.3 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 12 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.

10 Working with indexes (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

Users with the permission to use Index Manager can create new indexes, rebuild existing ones, and delete indexes. See Permissions for details on the Index Management permission.

Refer to Best practices for recommendations on preparing your data for indexing.

10.1 Navigating the Index Manager

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

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

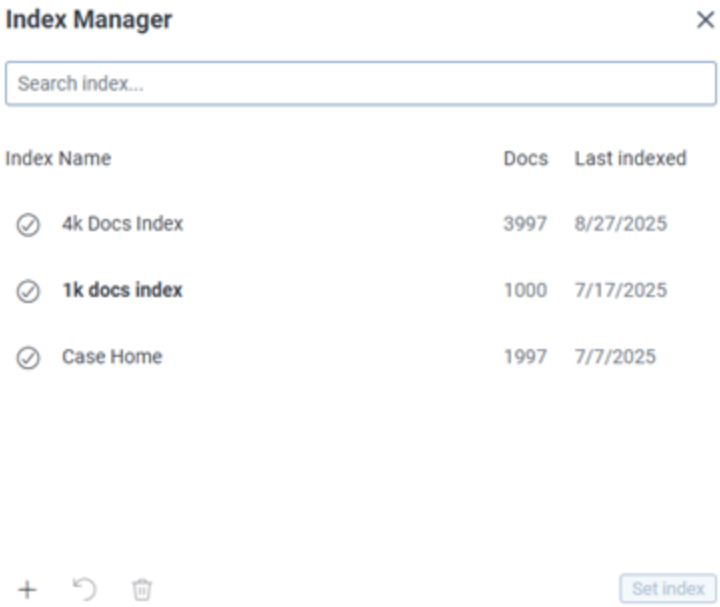

- **Search index box**—Use to search for a particular index in the list. Begin typing the name and matching indexes display for selection.
- **Index name**—name of the index entered during index set up. Hover over or click the index name to view index details, such as last indexed date/time and by whom, 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 23 for more information.
- **Delete index**—click to delete (trash can icon) the selected index. See Deleting an index on page 24 for more information.
- **Set index**—click button after selecting an index to use for querying.

10.2 Indexing rules and limitations

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

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

10.3 Creating a new index

Use Index Manager to create up to five indexes for different data sources (Saved Searches) per workspace.

To create a new index using Index Manager:

1. Access the aiR Assist chat panel by clicking its icon in the side bar.

2. From the chat panel, navigate to the Index Manager by clicking the index name in the **Ask a question** box.

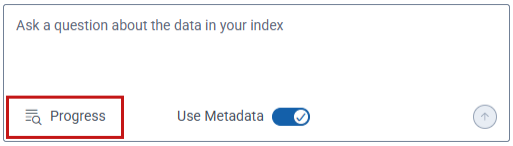

3. Click the **Create Index** (**+**) icon in the Index Manager list.

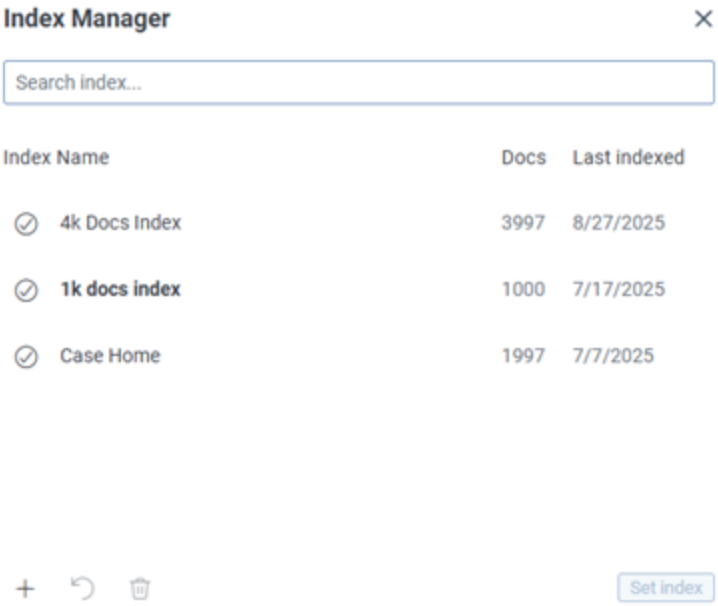

4. Fill in the following fields:

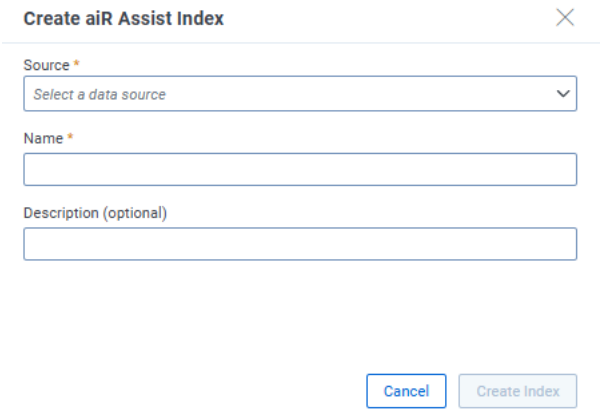

   - **Source**—choose a public saved search as the document source for your index. The index will be built from the extracted text of documents in that saved search. Only public saved searches are available for use.
   - **Name**—enter a unique, clearly descriptive name for the new index so its contents and purpose are identifiable by any user. Must be at least six 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: aiR Assist indexes are shared by all aiR 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 (Advanced Access) on page 20 for 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.
- aiR Assist indexes are shared by all users within the workspace.

10.4 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 aiR 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 (Advanced Access) on page 26 for more information.

Note: aiR for Case Strategy indexes cannot be rebuilt or deleted from Index Manager because they are managed by the aiR for Case Strategy application.

To rebuild an index from the Index Manager:

1. Access the aiR Assist chat panel by clicking its icon in the side bar.

2. Navigate to the Index Manager from the chat window by clicking the index name in the **Ask a question** box.

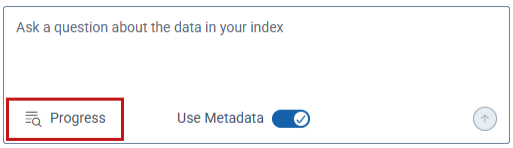

3. Select the desired index to rebuild.

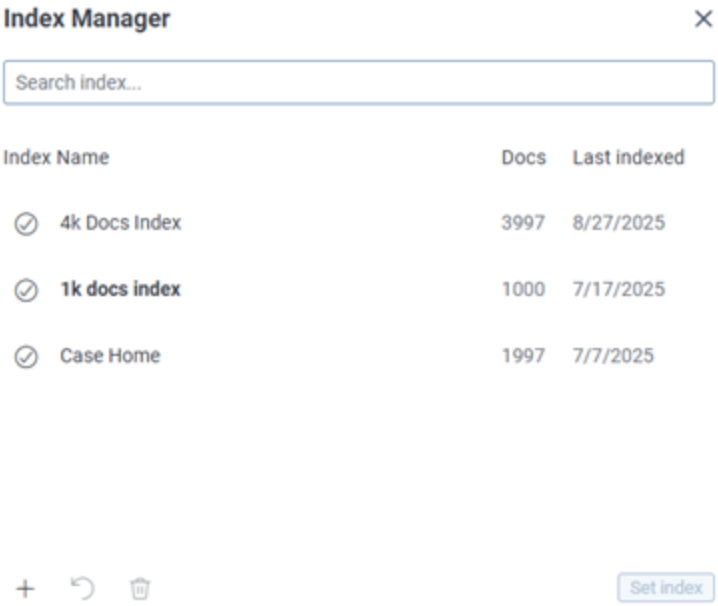

4. Click the **Rebuild** curved arrow icon. The index cannot be used during rebuilding, but you can work in other indexes or create a new one.
5. Click **Rebuild** to proceed.

Indexing progress appears below the new index name. See Monitoring the index build process on the next page for information.

10.5 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: aiR for Case Strategy indexes cannot be rebuilt or deleted from Index Manager because they are managed by the aiR for Case Strategy application.

To delete an index from the Index Manager:

1. Access the aiR Assist chat panel by clicking its icon in the side bar.

2. Navigate to the Index Manager from the chat window by clicking the index name in the **Ask a question** box.

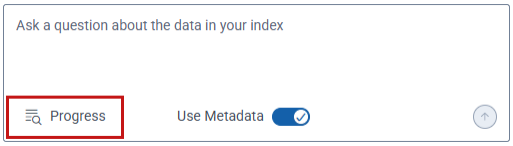

3. Select the desired index to delete.

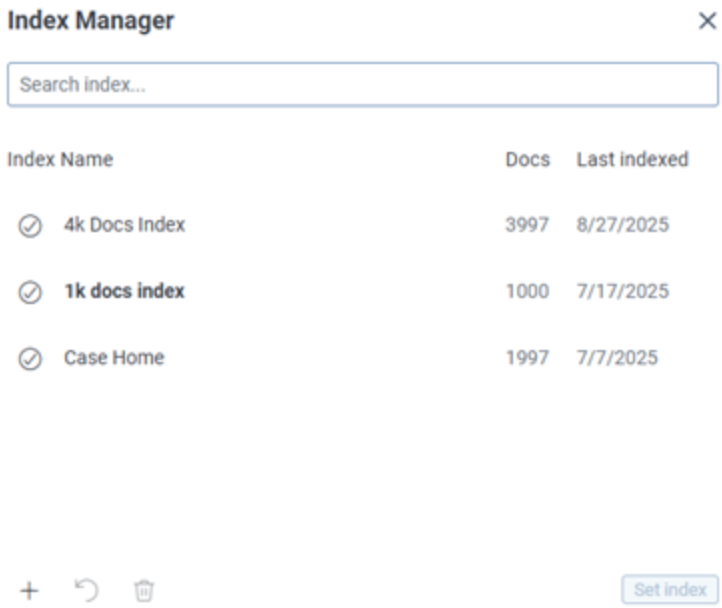

4. Click the **Delete Index** trash can icon.
5. Click **Delete** to proceed with permanently deleting the selected indexes. After the deletion process begins, it cannot be reversed.

10.6 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 aiR Assist.

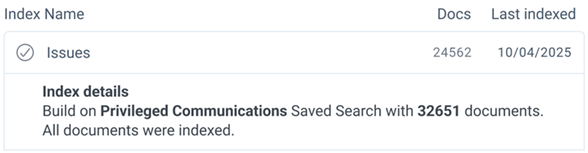

- **Partially built (Document errors encountered)**
  An exclamation mark within a triangle appears beside the Index Name to indicate a partial build, signifying that certain documents were not indexed. Index details show the name of the saved search used to create the index, total document count in 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 on the next page and Working with indexes (Advanced Access) on

page 20 for more information.

10.7 Error handling and retrying index jobs

aiR 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 (Advanced Access) on page 20.

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 aiR 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.

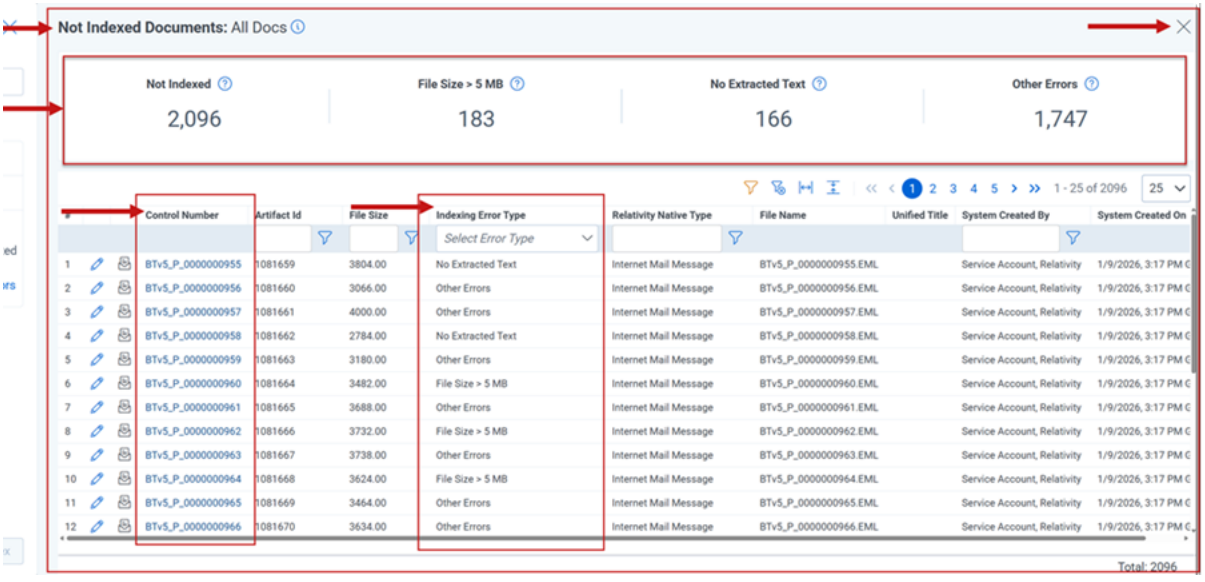

11 Metadata mapping (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality

described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

Incorporating metadata mapping helps improve how aiR 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, aiR 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, aiR 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.

aiR 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, aiR Assist helps to reduce noise and improves ranking quality. This approach is more precise than relying on semantic search alone. For more information on aiR Assist's usage of a Retrieval-Augmented Generation (RAG) architecture, see Understanding aiR Assist responses on page 4.

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 aiR Assist reasons over documents that are valid for the question being asked, reducing the likelihood of incorrect or over generalized responses.

Notes:

- Each workspace can have only one metadata mapping configuration.
- Metadata mapping is only supported for indexes created using aiR Assist. Currently, aiR for Case Strategy indexes do not support metadata mapping.

11.1 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 23 for more information.

1. Navigate to **aiR Assist**.
2. Select **Metadata Mapping** from the hamburger icon in the aiR Assist chat panel.

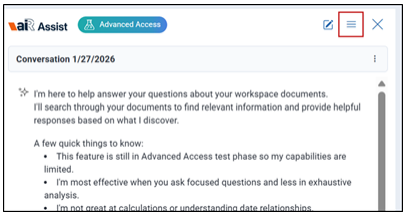

The Metadata Mapping panel slides open next to the aiR Assist chat panel. To close the panel, click the **X** opposite the panel heading or click **Metadata Mapping** again from the hamburger icon.

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

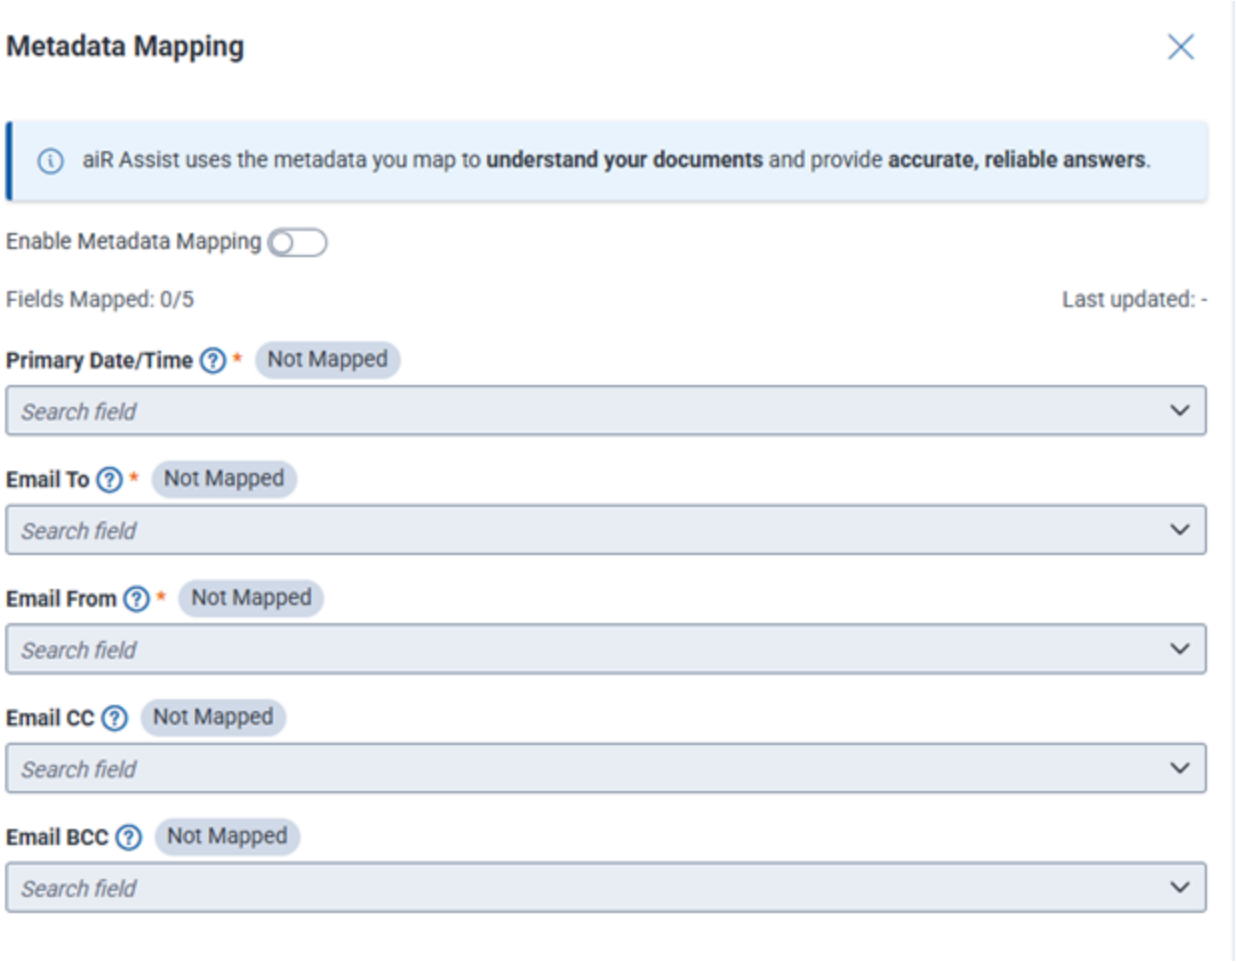

4. For each relevant metadata field listed that you want to use, select the appropriate workspace-level field from the corresponding drop-down. The fields listed in the drop-down will vary between workspaces. The first three metadata fields must be configured: **Primary Date/Time**, **Email to**, and **Email from**. Up to 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	Date	Required	Used to understand when document was created or sent in queries that require date filters.
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 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 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 and automatically updating all indexes.

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 25.

11.2 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.

Note: The Use Metadata toggle is only supported for indexes created using aiR Assist. Currently, aiR for Case Strategy indexes do not support metadata mapping so the toggle is disabled.

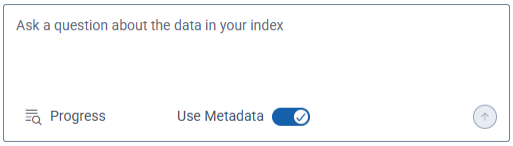

12 Frequently asked questions (Advanced Access)

Advanced Access (AA) is an opportunity to evaluate and work with Relativity features prior to the General Availability (GA) release. Relativity customers typically participate in AA programs on a feature-by-feature basis. The functionality described in this document may not be available in all Relativity environments. This document may not represent the functionality, appearance, or behavior of the GA release version of this feature.

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

12.1 Question FAQs

What types of questions can I ask in aiR Assist?
aiR Assist 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]?"

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.

How many references can aiR Assist return for a single question?
aiR Assist can return up to 25 references for one question, giving you a broad and well-supported set of sources.

Does aiR Assist ever reject my questions?
Yes. aiR Assist uses advanced intention capabilities to reject questions that cannot be answered or are outside its scope. It also requests clarification for vague or overly broad queries.

Are there recommended prompt patterns or techniques to improve accuracy for contract or relationship extraction tasks?
aiR Assist uses a Retrieval-Augmented Generation (RAG) approach to answer questions. 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 additional suggestions, see Best practices (Advanced Access) on page 9:

- Be explicit, not implicit
  Retrieval systems work best with clearly defined prompts. Vague questions such as “Tell me about the case”

are less effective than specific ones like “Describe contract terms related to indemnification.”

- Ask concise questions
  Multi-part or compound questions can fragment retrieval results. Breaking complex questions into smaller, focused prompts 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.

When using aiR Assist, it is also important to keep the following limitations in mind:

- Lack of comprehensiveness
  Retrieval systems aim to answer the question directly rather than provide a full list of all related documents or entities.
- No metadata support
  Queries that depend on document metadata (such as, file type, custodian, or date range) may return incomplete or imprecise results, as metadata filtering is not currently supported.

Can aiR Assist understand multiple languages?
aiR Assist is designed and optimized for English-language content. It has not been formally tested or validated for use with multilingual or non-English datasets. While it may process and return responses for non-English text, accuracy and completeness cannot be guaranteed, and verification of cited sources is strongly recommended.

Can I share individual questions asked in Air Assist with others?
No. You cannot share individual questions with others within aiR Assist.

How can I limit users to only asking questions, without permission to create, rebuild, or delete indexes?
You must set up their permissions to only allow access to Prompting. See Permissions for details.

12.2 Answer/Response FAQs

Can aiR Assist be configured to produce more consistent or deterministic answers?
No. aiR Assist does not currently include configuration options to control response variability or “thinking time.” The system prompt used by aiR 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 question, the more consistent the output tends to be. aiR Assist focuses on accuracy and relevance rather than adjustable response styles. See LLM model in use on page 4.

Does aiR Assist filter out or de-prioritize repeated or boilerplate content (such as disclaimers or email signatures)?
Not at this time. aiR 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?
aiR Assist 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. Using the feedback buttons in the response box can help improve the tool over time.

How can I leave feedback (like or dislike) within the aiR Assist application?
To leave feedback in the aiR Assist app, use the thumbs up or down icons below each answer to rate it and add comments. See Prompting (Advanced Access) on page 10 for more information. You may also occasionally be present with an aiR Assist Feedback dialog asking how helpful were aiR Assist's answers for completing your task. See Prompting (Advanced Access) on page 10 for details.

12.3 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.

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 history persists between your RelativityOne user sessions, unless you delete the conversation. See Using Conversation Manager for details on deleting conversations.

Can I share conversations with other users in Air Assist?
No. You cannot share conversations with others within aiR Assist.

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 ask more questions?
You can switch between conversations anytime, as well as return to a previous conversation to continue asking questions using Conversation Manager. See Using Conversation Manager for details.

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

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

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

12.4 Index FAQs

Where are aiR Assist indexes stored, and how can indexing errors be reviewed?
aiR Assist indexes do not currently have a dedicated view within RelativityOne. Indexes are managed directly through the aiR 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 are automatically skipped during the indexing process and are not currently displayed as errored items.

Are there any known delays or issues that cause aiR Assist to ignore newly indexed documents?
There are no known delays or edge cases that prevent aiR Assist from accessing newly indexed documents. Once the index build task is complete, all included documents become immediately available for querying. Relativity offers workflows that can identify repeated text.

How can I limit users to only manage indexes (create, rebuild, or delete), without permission to ask questions?
You must set up their permissions to only allow access to Index Management. See Permissions for details.

12.5 Document FAQs

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

aiR Assist generates its answers based solely on the extracted text from processed documents. This means that for Excel files, the content available to aiR 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 aiR Assist able to reliably read documents that contain only images and no native text files?
aiR Assist can only use the extracted text available in RelativityOne. 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 RelativityOne.

Without extracted text, aiR Assist will not be able to index or generate answers 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 size limit. During indexing, aiR Assist automatically divides large documents into smaller text chunks to optimize retrieval and improve response relevance. These chunks may be retrieved individually when aiR Assist processes a query, ensuring that large documents are handled efficiently without additional configuration.

Are documents larger than 5 MB flagged in the error report?
No. Documents exceeding the 5 MB 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 aiR 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.

13 Index