exterro®
Data Risk Management

FTK ASSIST FEATURE GUIDE

JULY 2026

Table of Contents

About Exterro

Exterro was founded with the simple vision that applying the concepts of process optimization and data science to how companies manage digital information and respond to litigation would drive more successful outcomes at a lower cost. We remain committed to this vision today. We deliver a fully integrated Data Risk Management platform that enables our clients to address their privacy, regulatory, compliance, digital forensics, and litigation risks more effectively and at lower costs. We provide software solutions that help some of the world's largest organizations, law enforcement and government agencies work smarter, more efficiently, and support the Rule of Law.

1 Introduction to FTK Assist

Investigations generate vast amounts of evidence spanning thousands of documents, emails, and chat conversations that investigators must review to build understanding and draw conclusions. Reading through this volume of material manually is time-consuming and prone to oversight. Critical information can be buried across dozens of files, and connecting the dots between dispersed pieces of evidence demands sustained effort that slows investigations down.

FTK Assist is an AI-powered conversational feature built directly into FTK Central that transforms how investigators interact with case evidence. Rather than manually searching through files, investigators can simply ask questions in plain English and receive precise, evidence-backed answers in seconds. Every answer is grounded exclusively in the processed artifacts of the case — documents, emails, and chat conversations, and includes citations that trace each response back to its source.

Note: *FTK Assist is not a general-purpose chatbot. It does not draw on external knowledge or make assumptions beyond what the evidence contains. It is a focused investigative tool designed to surface relevant information from within a case, accelerate discovery, and support more confident, evidence-backed decisions.*

2 Key Capabilities

FTK Assist brings together several capabilities that collectively change the pace and depth of an investigation.

Capability	Description
Natural Language Queries	Ask questions in plain language — no complex search syntax or Boolean operators required.
Source-Grounded Answers	Every response is generated strictly from processed case artifacts. FTK Assist never fabricates information or draws on external knowledge.
Inline Citations	Each answer includes citations linked to the specific source document, enabling immediate verification.
Conversational Context	Follow-up questions are supported within the same conversation, allowing investigators to progressively explore a topic without starting over.
Apply as Filter	A single click applies the artifacts referenced in any answer as a filter in the Review grid, enabling immediate deeper review of relevant evidence.
Case Isolation	FTK Assist is strictly case-scoped. Answers are generated only from the artifacts of the current case, preventing cross-case data leakage.
Conversation History	Up to 50 previous conversations within a case are preserved per user, allowing investigators to revisit earlier questions and answers.

3 How FTK Assist Works

FTK Assist operates in two stages: preprocessing and query execution. Understanding this workflow helps investigators make the most of the feature.

3.1 Preprocessing

Before investigators can ask questions, supported case artifacts such as documents, emails, and chat conversations must be preprocessed for FTK Assist. During preprocessing, each artifact is:

- Split into smaller semantic chunks that preserve contextual meaning
- Converted into vector embeddings that capture the semantic meaning of the content using the **Qwen3-Embedding-0.6B model** (1024 vector dimensions)
- Stored in a centralized PostgreSQL vector database for efficient retrieval

Artifacts are batched and dispatched to available AI Server nodes in parallel. In environments with multiple AI Server nodes, a load balancer distributes the workload across nodes, reducing preprocessing time for large case datasets.

Note: *Only artifacts that have been preprocessed for FTK Assist are included in the conversational scope. Artifacts that have not been processed are excluded from all answers.*

3.2 Query Execution

When an investigator asks a question, FTK Assist automatically classifies the query and executes the appropriate retrieval strategy (as given below). There are three query types, and the system determines which applies without needing the user to specify.

Query Type	How It Works	Example Question
Semantic Search	The query is converted into an embedding and compared against the vector database. Semantically relevant chunks are retrieved and passed to the LLM to generate a contextual answer. Answers include inline citations.	*"What was discussed about Project X in the communications?"*
Metadata Search *(Beta)*	The query is interpreted as a structured data request. An SQL query is generated and executed against case metadata, scoped to emails and chat conversations. The result is returned as a natural language answer.	*"How many emails were shared between A and B?"*
Hybrid Search *(Beta)*	A combination of both approaches. Used when a question requires both metadata retrieval and contextual understanding. The system generates one unified response with citations.	*"Summarize all email conversations between A and B."*

Note: *Metadata search is scoped to emails and chat conversations only. Documents are not included in the metadata search scope, i.e., structured metadata-based questions such as counts cannot be answered for documents.*

4 Enabling FTK Assist

Before investigators can use FTK Assist, supported artifacts such as documents, emails, and chat conversations must be preprocessed. During preprocessing, each artifact is broken down into smaller, meaningful pieces through a process called chunking. These chunks are then converted into vector embeddings, and stored in a vector database. When an investigator asks a question, relevant chunks are retrieved from this database and used to generate a precise, evidence-backed answer.

Preprocessing can be initiated through two workflows in FTK Central: during evidence processing, or after evidence has already been processed via Additional Analysis.

4.1 During Evidence Processing

This is the recommended approach when adding new evidence to a case. Enabling FTK Assist at this stage ensures that supported artifacts are preprocessed for Assist as part of the same evidence processing job, with no additional steps required afterward.

1. After creating the case, navigate to the **Process Evidence** section.
2. Click **Customize Options** at the bottom of the page.
3. In the customization window, select **AI Jobs** from the left menu.

4. Select **Data Assist** from the right pane.

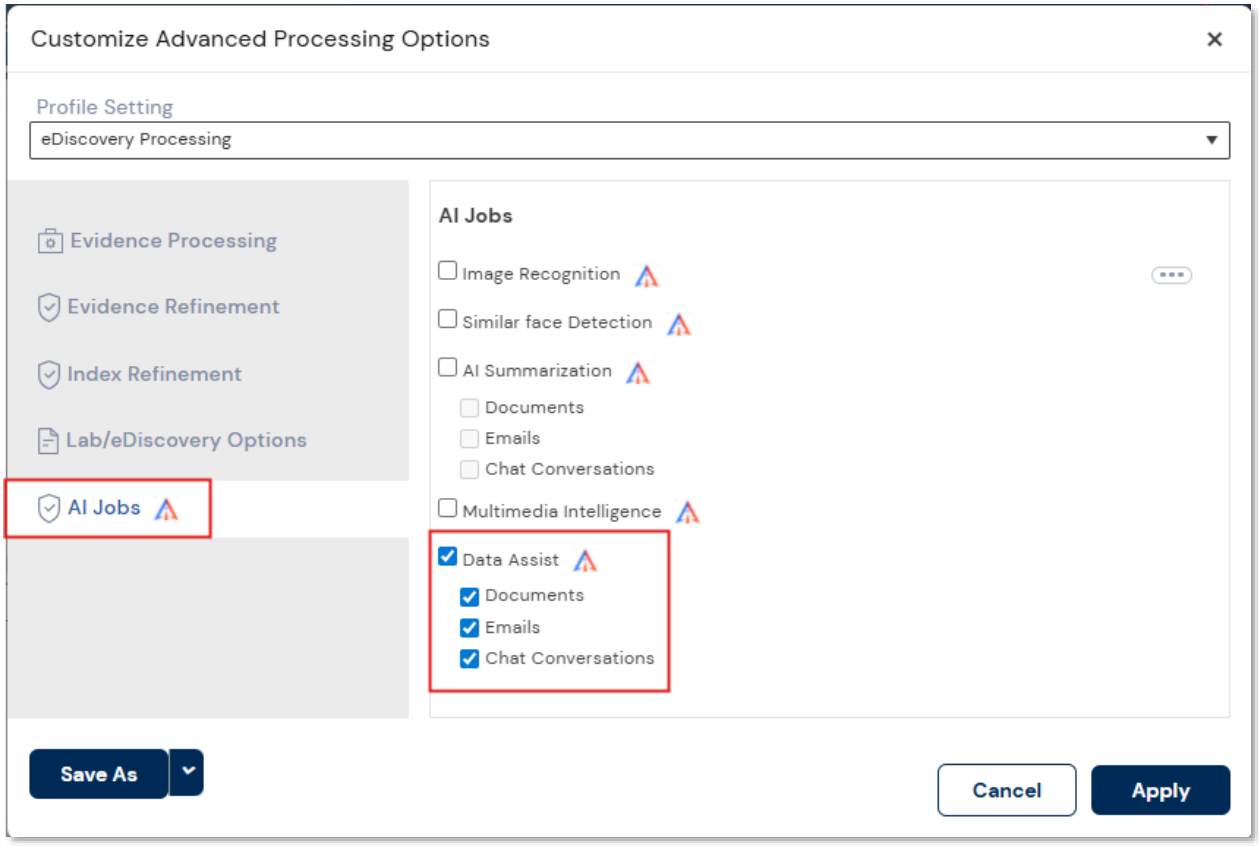

5. Click **Apply**, and then click **Process Data** to begin evidence processing.

Documents, emails, and chat conversations in the evidence will be preprocessed for FTK Assist alongside the standard evidence processing job, according to the artifact selections made.

4.2 After Evidence Has Been Processed

If evidence has already been processed and FTK Assist preprocessing was not enabled at that time, it can be initiated for specific artifacts using the Additional Analysis workflow. This is also useful when new artifacts are added to an existing case and need to be brought into the FTK Assist scope.

1. Navigate to the **Home** tab of the Review page.
2. Select the artifacts that need to be preprocessed for FTK Assist.
3. Right-click on the selected artifacts and choose **Additional Analysis**.

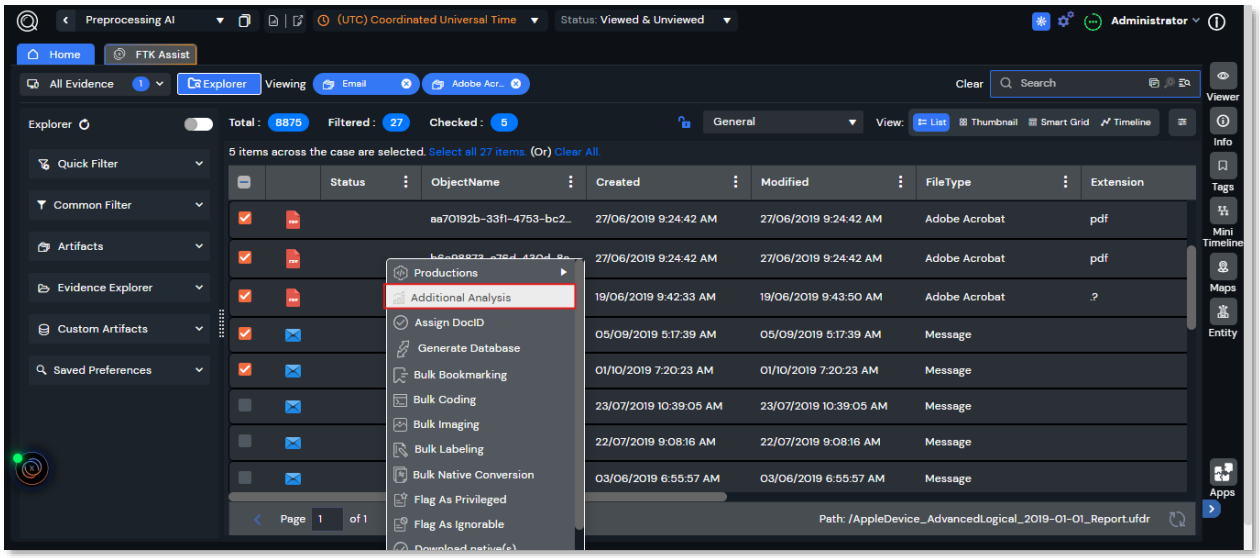

4. In the Additional Analysis window, click **Customize Options**.

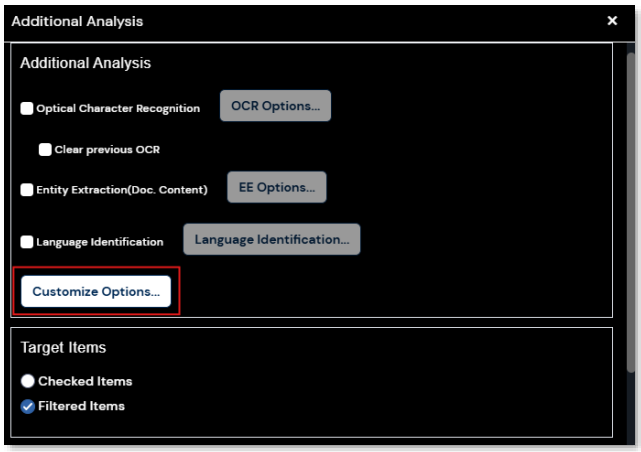

5. Select **AI Jobs** from the left menu.
6. Select **Data Assist** from the right pane. If needed, check or uncheck **Documents**, **Emails**, and **Chat Conversations** individually based on your requirements.

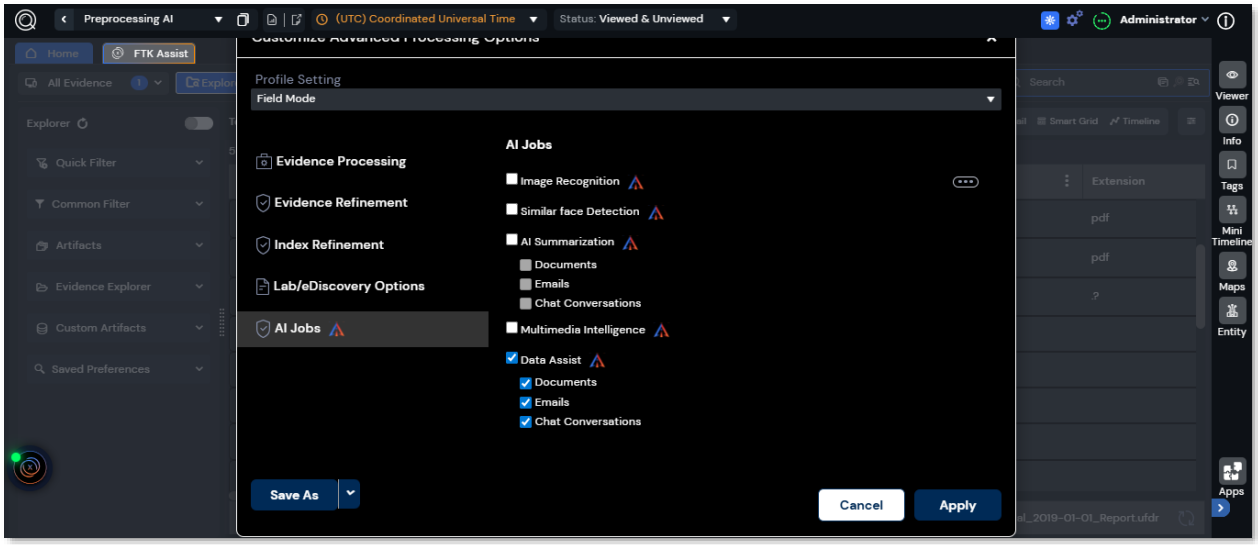

7. Click **Apply**.

The selected documents, emails, and chat conversations will be queued for FTK Assist preprocessing.

Note: *If new artifacts are added to a case after initial preprocessing, they will not automatically be included in the FTK Assist conversational scope. These artifacts must be preprocessed manually through the Additional Analysis workflow before FTK Assist can use them to generate answers.*

4.3 Monitoring Preprocessing Progress

Once preprocessing is initiated through either workflow, a job is created in the FTK Central Job Queue. The preprocessing job appears as **AI - Assist** and can be monitored for progress and status from the Job Queue.

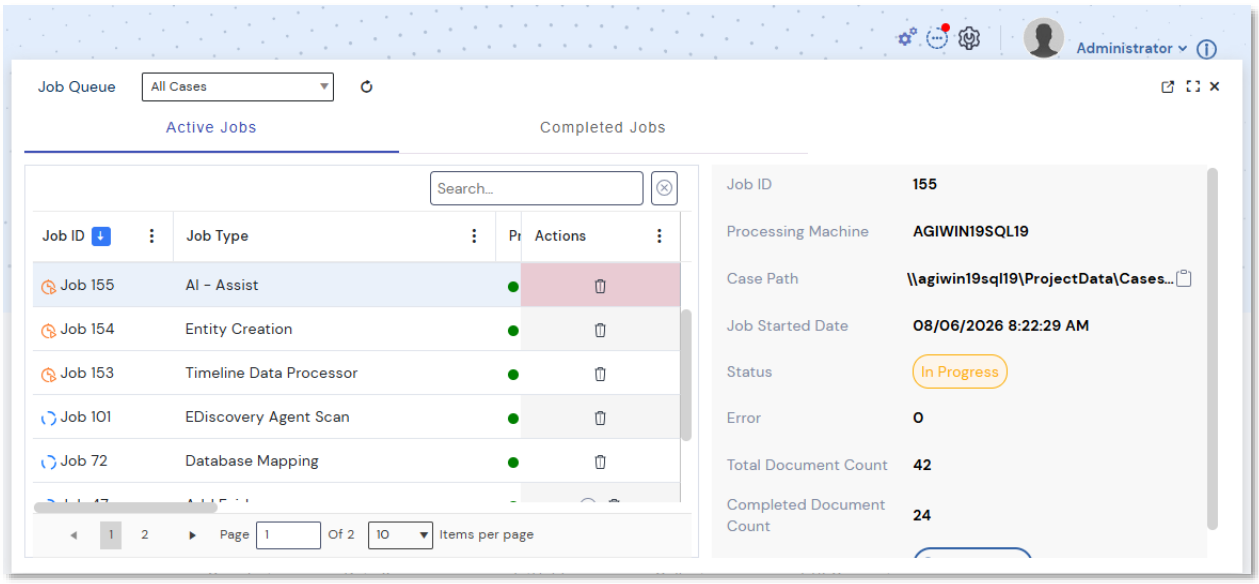

4.4 Filtering Preprocessed Artifacts

To quickly identify which artifacts have been preprocessed for FTK Assist and which have not, use the Processed for Assist filter available under Quick Filters in the Review grid.

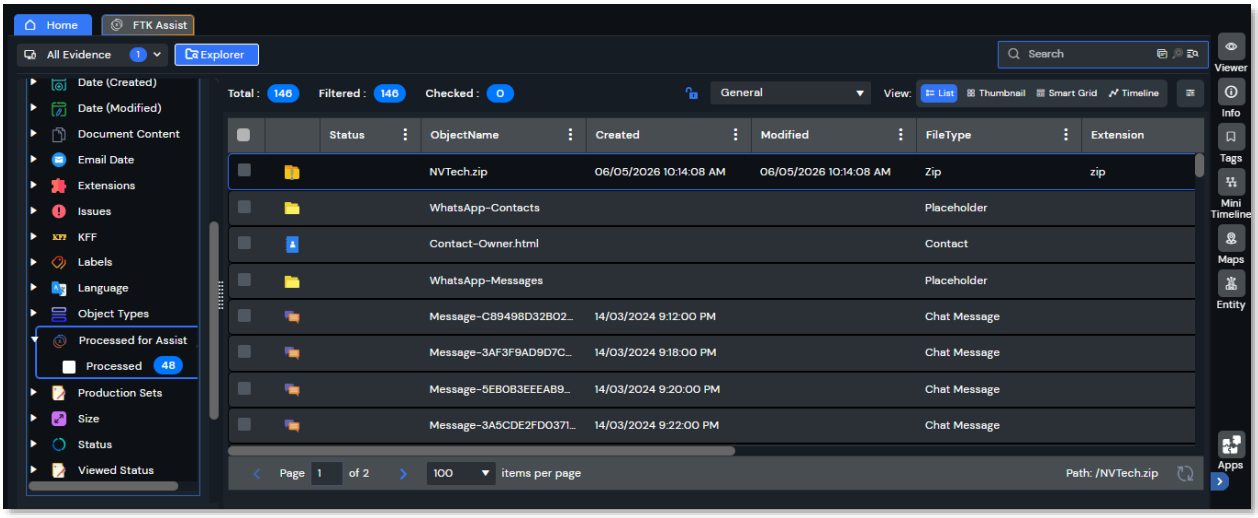

The filter provides two options:

- **Processed** — Lists all artifacts that have already been preprocessed for FTK Assist and are within the conversational scope.
- **Not Processed** — Lists artifacts that are yet to be preprocessed and will not be included in FTK Assist answers until they are.

This filter is particularly useful before running Additional Analysis, as it allows investigators to quickly identify and select only the artifacts that still need to be processed, without having to manually scan through the full evidence set.

Note: *The 'Processed for Assist' filter is scoped exclusively to the artifact types supported by FTK Assist, i.e., documents, emails, and chat conversations. Other artifact types present in the case are not included in the filter counts. The Processed and Not Processed numbers reflect only the distribution of preprocessing status within the supported artifact types, and not the total number of objects in the case.*

5 Using FTK Assist

5.1 The FTK Assist Tab

FTK Assist is available as a dedicated tab in FTK Central, located directly next to the Home tab. Once artifacts have been preprocessed, investigators can open the Assist tab and begin asking questions immediately.

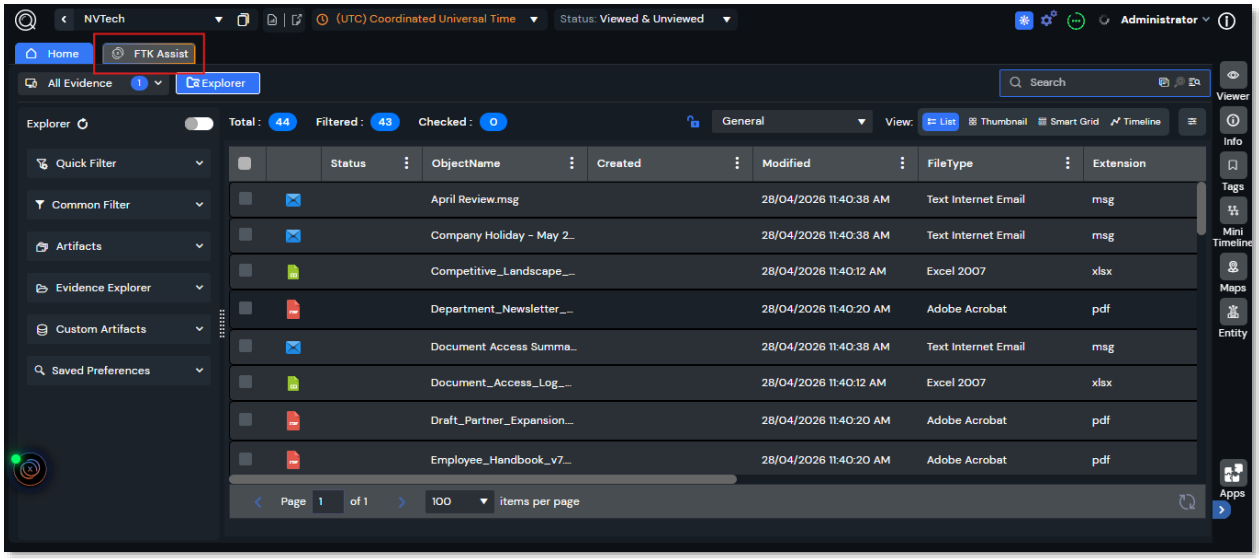

An Info button located in the upper-right section of the FTK Assist interface displays the count of artifacts that have been preprocessed and are currently in scope for analysis within FTK Assist. This gives investigators a clear picture of what evidence base their questions are being answered from before they begin.

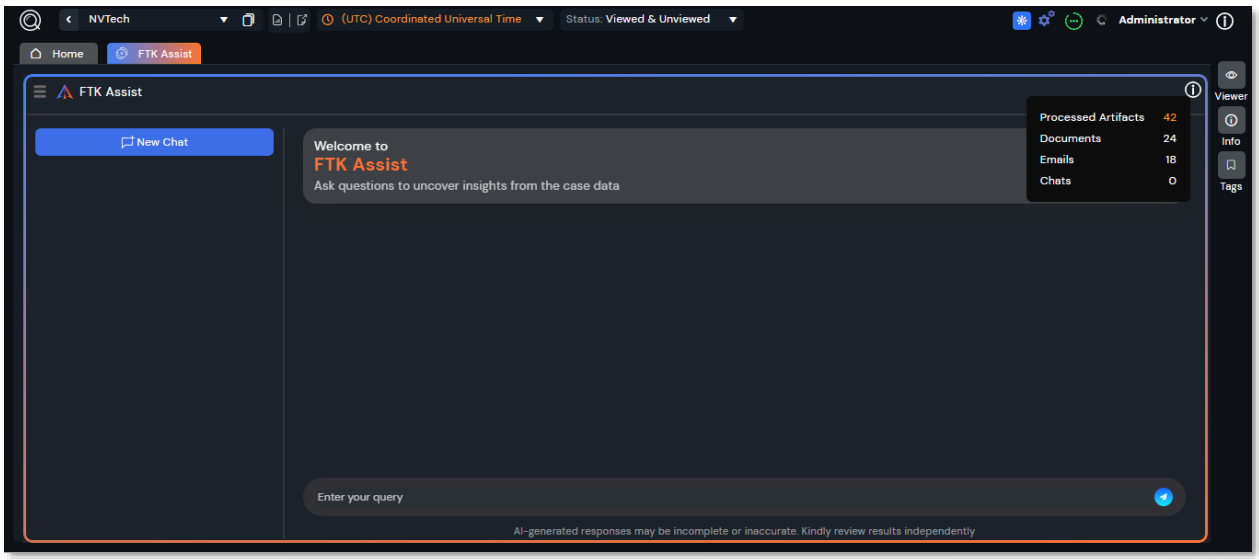

5.2 Asking Questions

To ask a question, type it into the chat bar at the bottom of the Assist window and submit.

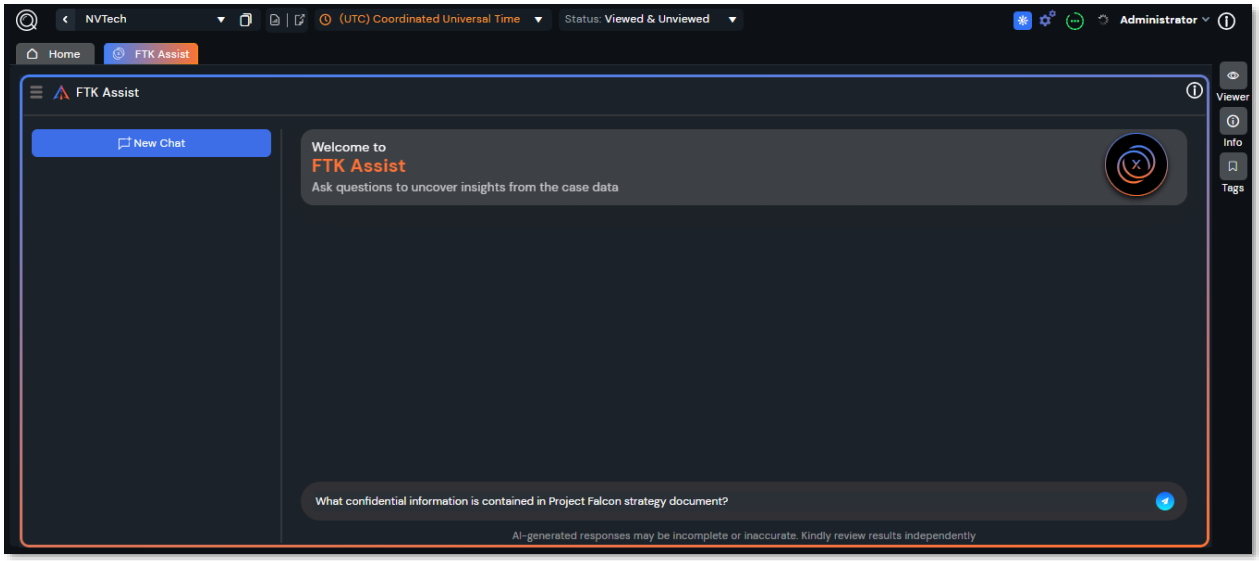

Note: *For guidance on querying FTK Assist effectively and best practices for obtaining relevant results, refer to* ***Appendix A – Prompting Best Practices***.

5.3 Traces and Answer Generation

As soon as a question is submitted, the **Traces** section appears in real time, showing the step-by-step process being executed in the background, from query classification and chunk retrieval through to answer generation. Investigators can watch this unfold as it happens, giving full visibility into how the system is working to arrive at an answer.

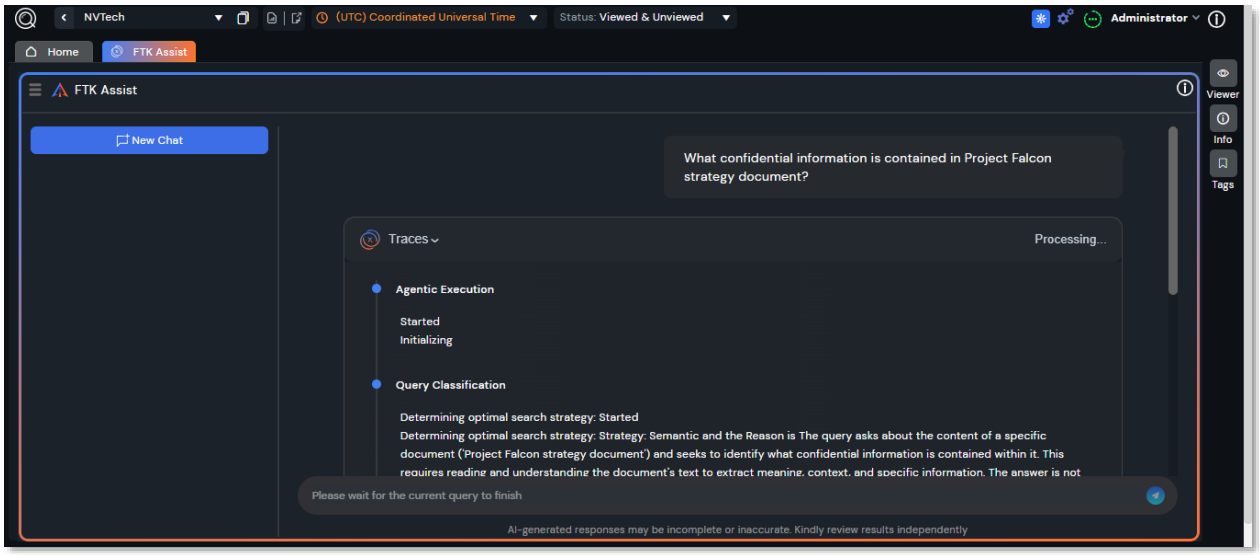

Once processing is complete, the answer is displayed inline within the conversation, directly below the Traces. Investigators can continue to view the Traces even after the answer has been generated, allowing them to revisit the reasoning and retrieval steps behind any response at any time. This is designed to ensure full explainability of the AI every answer comes with a transparent record of how it was produced.

Follow-up questions are supported and answered using the existing conversational context, allowing investigators to progressively explore a topic within a single conversation without needing to re-establish context each time.

Note: *The Qwen3-30B-Instruct-2507 model generates answers using only the retrieved case artifacts as input. If no relevant content is found in the processed artifacts, FTK Assist explicitly states that it does not have sufficient context to answer the question rather than speculating or fabricating a response. If the AI Server or LLM Server is unreachable, FTK Assist will display an error message in the interface. If this occurs, contact your system administrator to verify that the servers are running and accessible.*

5.4 Answers and Inline Citations

Answers are displayed inline within the conversation. For semantic and hybrid queries, each statement or claim in the response includes an inline citation — a reference marker that links that specific part of the answer back to the exact source artifact from which it was derived. This ensures that every assertion in a response can be traced directly to the underlying evidence, allowing investigators to verify the answer rather than take it on trust.

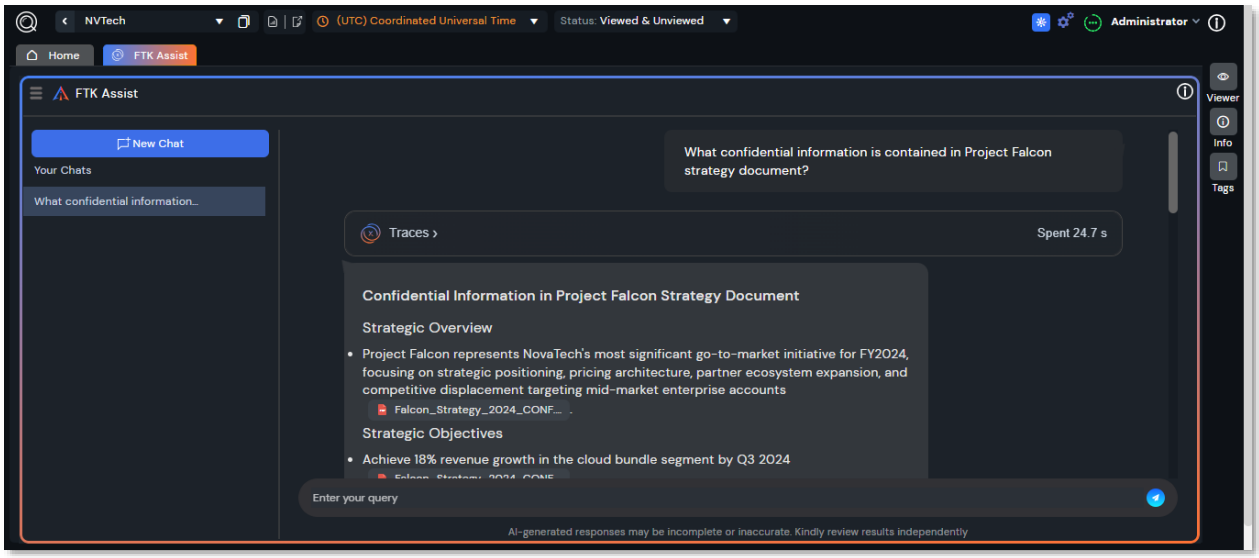

5.4.1 Viewing the Source Document

Clicking on an inline citation opens the referenced artifact in its entirety in a viewer pane on the right side of the Assist tab. The document can be viewed using the same native, image, and text viewing options available elsewhere in FTK Central.

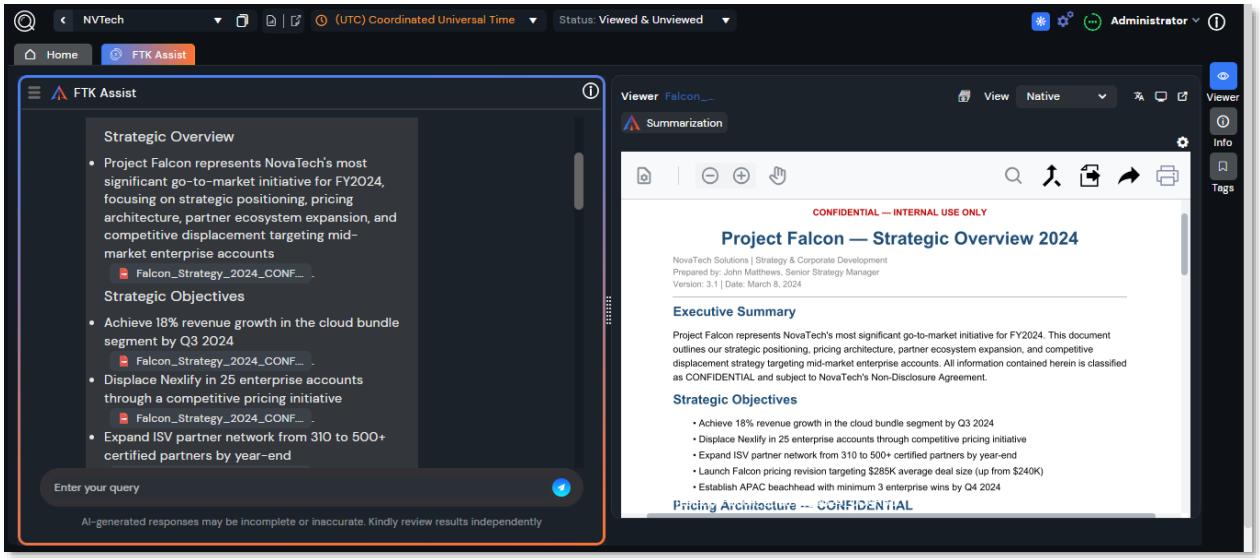

5.4.2 Info and Bookmark Panes

Within the Assist tab, investigators can also access the **Info pane** for the opened document, which displays artifact metadata and details. The **Bookmark pane** allows investigators to add labels and bookmarks to the artifact directly from within the Assist tab, without needing to navigate back to the Review grid.

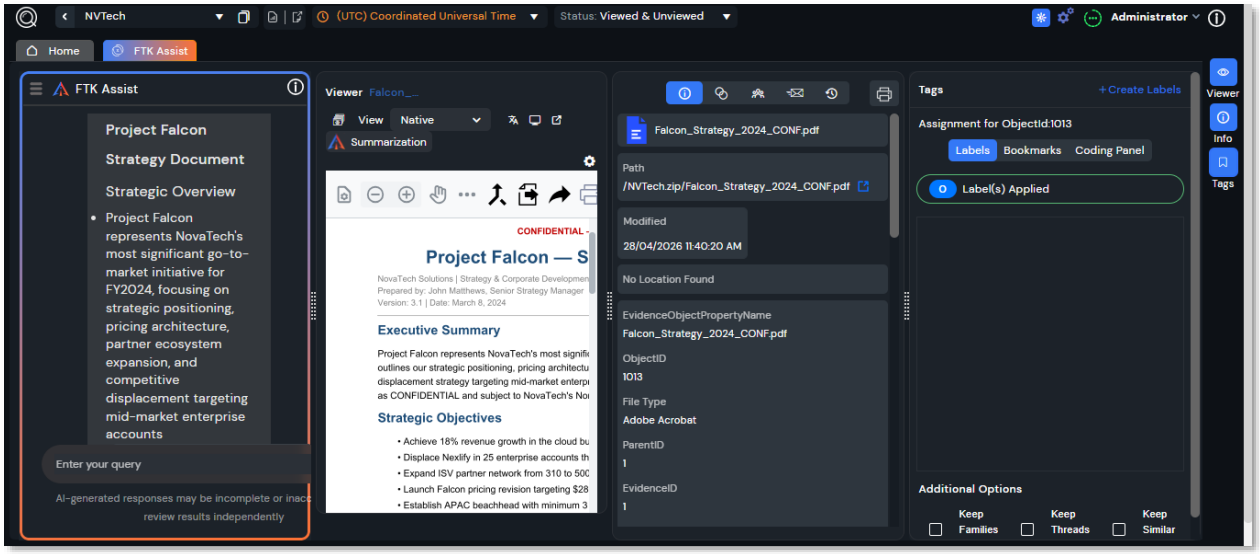

5.5 Sources and Bulk Actions

At the bottom of each answer, a **Sources** dropdown lists all the artifacts that were used to generate the response. Investigators can expand this dropdown to see the full list of contributing source documents.

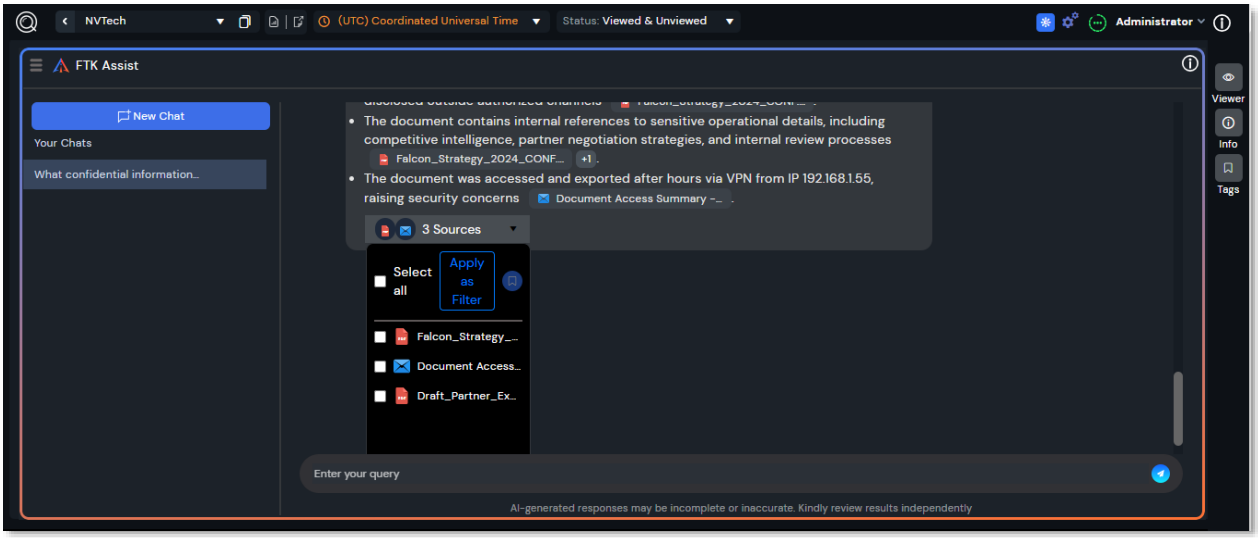

From the Sources dropdown, investigators can select all or a subset of the listed artifacts and perform bulk labeling directly by tagging multiple relevant artifacts in a single action without leaving the Assist tab.

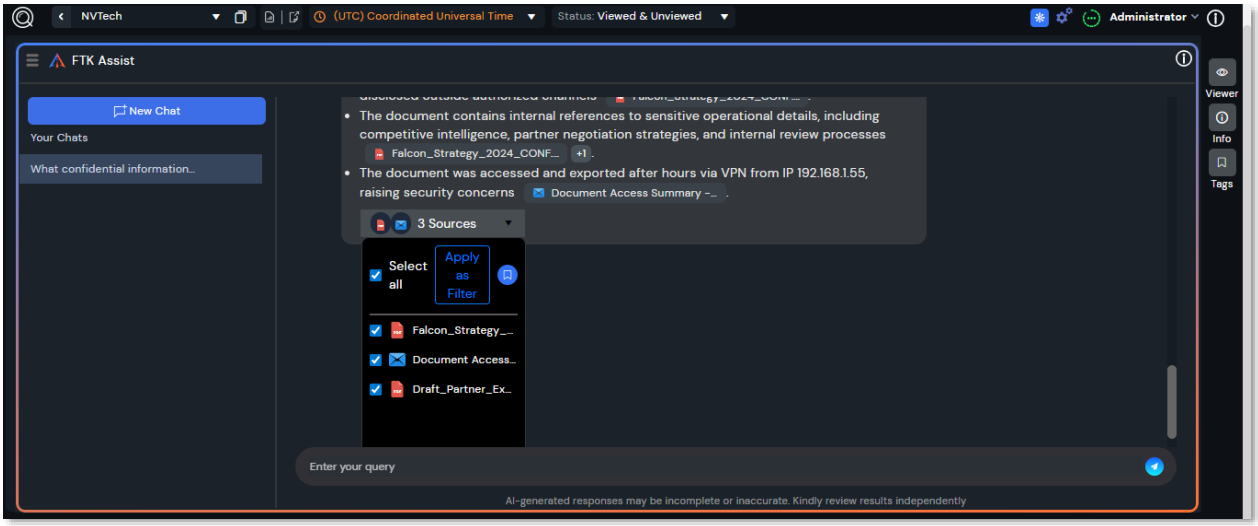

Additionally, an ‘Apply as Filter’ button is available. Clicking this applies a filter on the Home tab that displays only the artifacts selected from the Sources dropdown, allowing investigators to pivot immediately from an FTK Assist response into a focused review of the relevant evidence in the full Review interface.

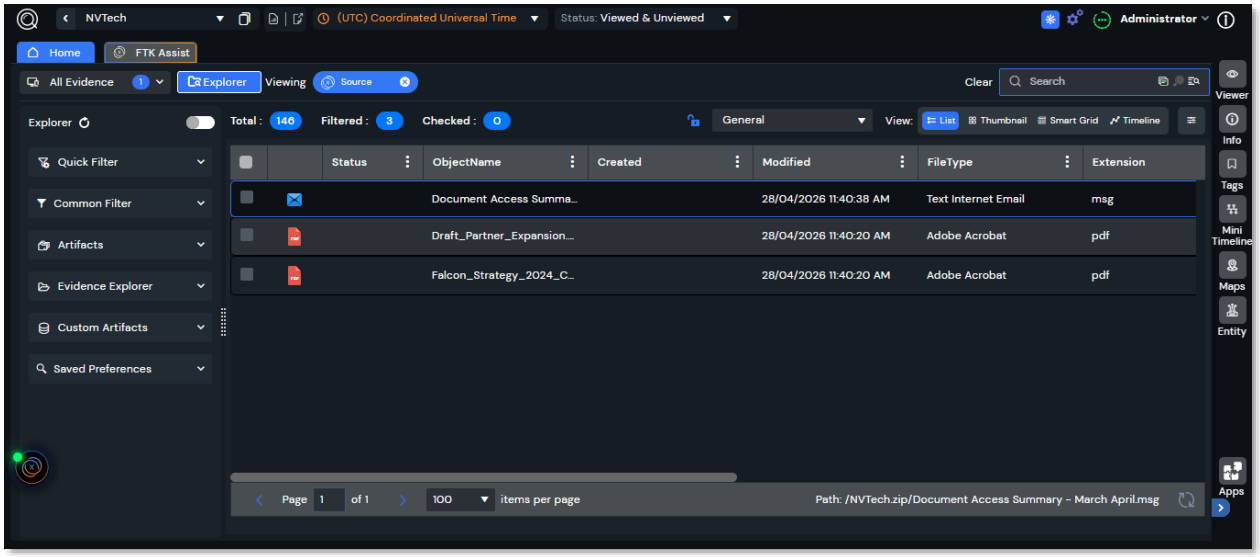

Note: *Metadata search is a beta capability in this release. For metadata-based queries where the answer is derived from structured case metadata rather than document content, inline citations, the Sources dropdown, Apply as Filter, and Label are not available.*

5.6 Conversation History

All conversations within FTK Assist are stored at the case level and user level. This means each conversation is visible only to the user who initiated it, and only within the case in which it was created. It will not be visible to any other user, nor will it appear under any other case.

To access previous conversations, click the hamburger menu icon at the top left of the Assist window. This menu displays all past conversations for the current user within the current case, and also contains the button to start a new conversation.

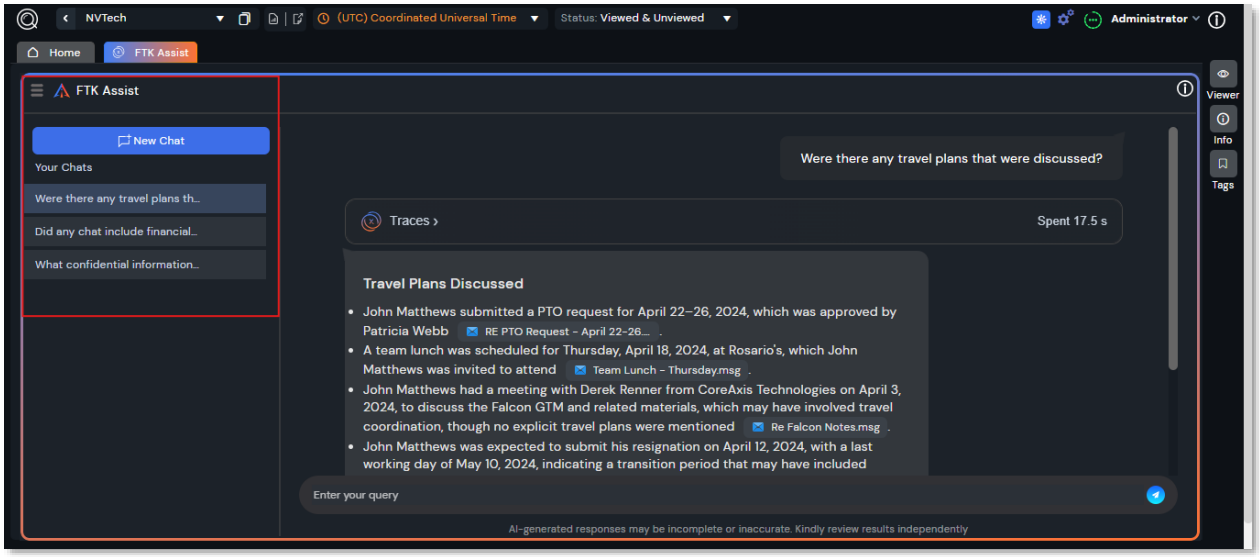

FTK Central stores the most recent 50 conversations per user per case. There is no limit on the number of questions that can be asked within a single conversation — the full question and answer history of each conversation is preserved in its entirety.

For deleting conversations, users require the appropriate delete permission, which is configured in User Management. Refer to the section – *User Permissions*.

5.7 Floater Window

In addition to the dedicated Assist tab, FTK Assist is also accessible as a floating window from the Home tab. The Assist icon at the bottom left of the application opens a floater window that provides the same full FTK Assist functionality, allowing investigators to ask questions and review answers while simultaneously browsing the Review grid or other parts of the interface without switching tabs.

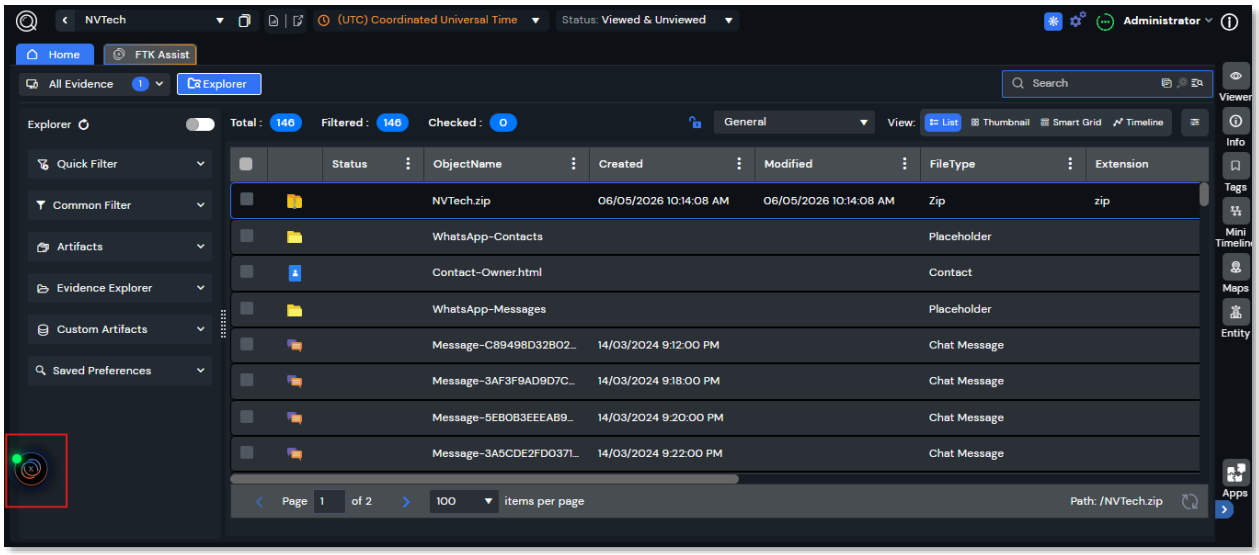

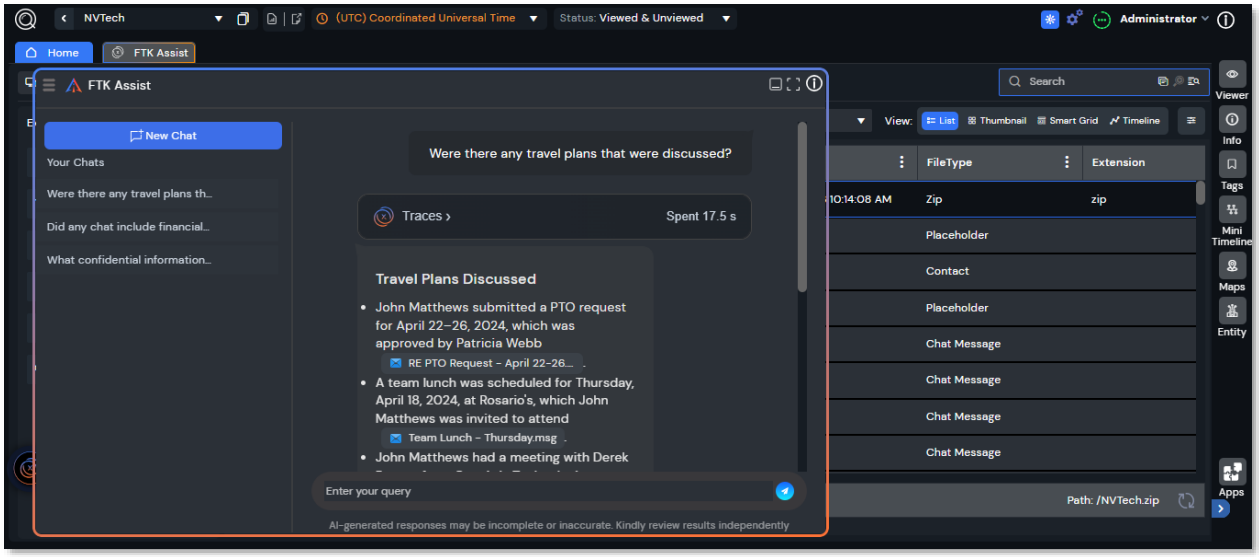

The floater window icon can be hidden if it is not needed. Once hidden, the icon will not reappear until the page is refreshed or the user logs in again.

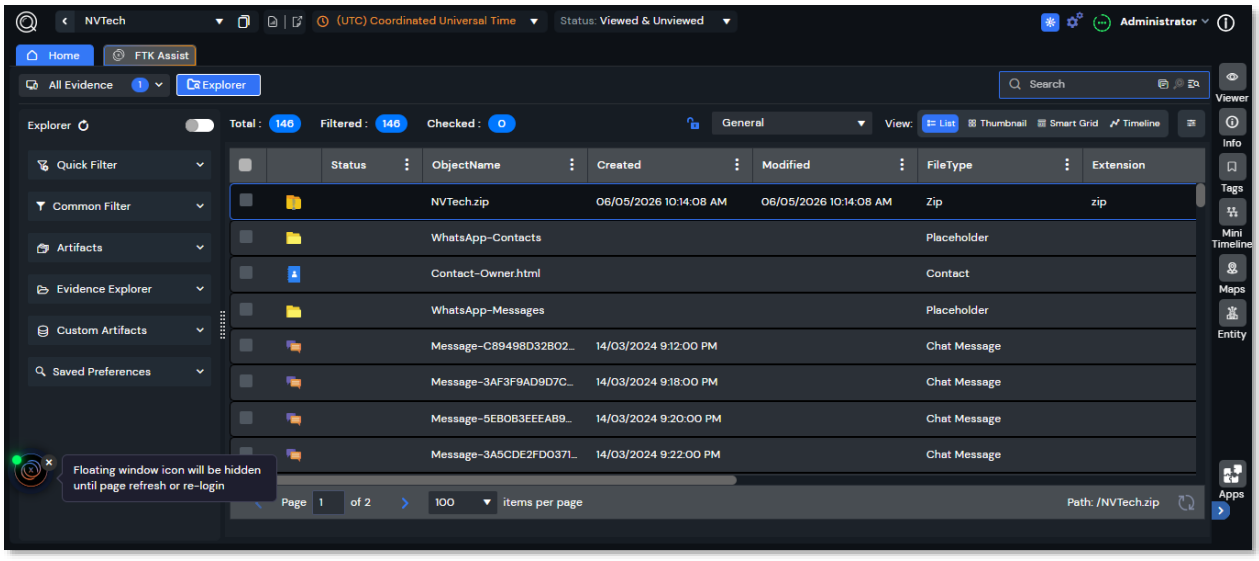

Note: *FTK Assist respects user-level access permissions. For users who do not have access to privileged artifacts, those artifacts are excluded from the retrieval scope and will not be used to generate answers. This ensures that FTK Assist responses are always consistent with the access rights of the investigator asking the question.*

6 User Permissions

Access to FTK Assist is controlled through user permissions configured in the 'User Management' section of FTK Central Settings. Permissions are applied globally: if a user has been granted access to FTK Assist, they can use it across all cases they have access to.

There are two configurable permission levels for FTK Assist:

Permission	Description
Access	Allows the user to open FTK Assist, ask questions, and use all features within the Assist interface. Users without this permission cannot access FTK Assist.
Delete	Allows the user to delete conversations within FTK Assist. Users with the Access permission but without the Delete permission can view conversations but cannot delete them.

Permission configuration is performed by administrators via 'Administration' > 'User Management'. In the User Management section, navigate to 'Role Permission Mapping'. Click 'Edit' against the role for which permissions need to be configured. Under the role's permissions, locate 'Data Assist' and check or uncheck it and the permissions listed under it — Access and Delete, as required.

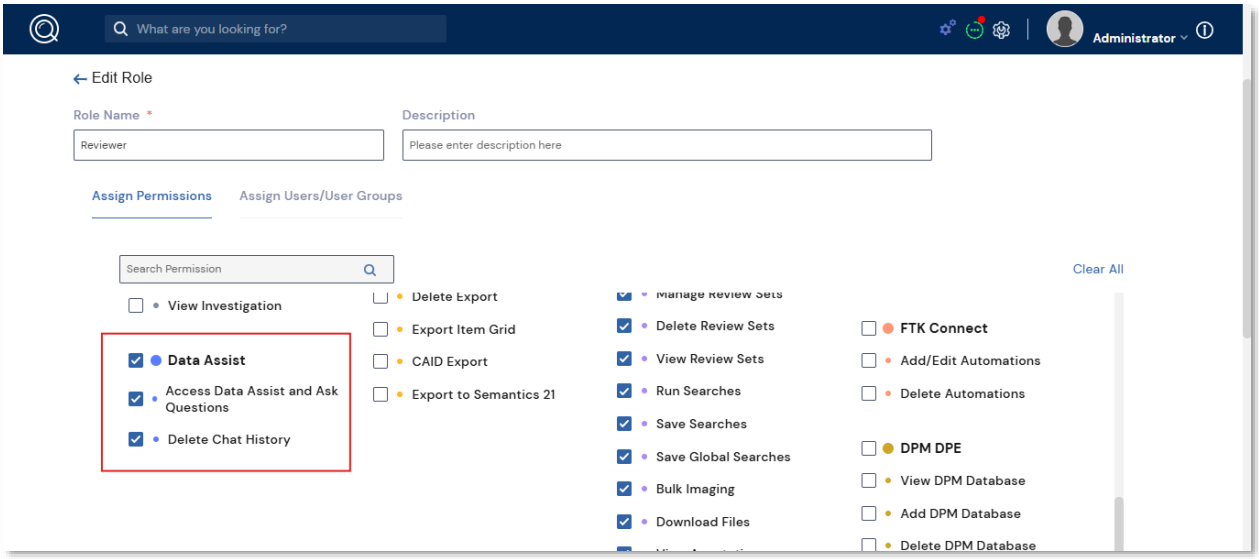

7 Important Notes

- FTK Assist only answers questions based on artifacts that have been preprocessed for Assist. Artifacts that have not been processed are excluded from the conversational scope.
- If preprocessing is still in progress when a query is submitted, answers will be generated from the artifacts that have already been fully processed. If no artifacts have completed preprocessing, users must wait for preprocessing to complete before submitting queries.
- If no relevant information is found in the processed artifacts, FTK Assist will explicitly state that it does not have enough context to answer the question. It will not speculate or fabricate a response.
- The query classification (semantic, metadata, or hybrid) is determined automatically by the system. Users do not need to specify the query type.
- Metadata search is a beta capability in this release. Inline citations, the Sources dropdown, Apply as Filter, and Label are available for semantic and hybrid queries; metadata search responses are presented as natural language answers only.
- Metadata search is scoped to emails and chat conversations only. Documents are not included in the metadata search scope, meaning structured metadata-based questions (such as counts) cannot be answered for documents.
- Conversation history is scoped per user, per case. One investigator cannot view the conversation history of another investigator within the same case.
- A maximum of 50 conversations are stored per user per case. All questions and answers within each conversation are preserved in full.
- FTK Assist is strictly case-scoped. It does not access artifacts from other cases and does not draw on any external knowledge sources.
- FTK Assist respects user-level access permissions. For users who do not have access to privileged artifacts, those artifacts are excluded from the retrieval scope and will not be used to generate answers. Responses are always consistent with the access rights of the investigator asking the question.

8 Appendix A – Prompting Best Practices

FTK Assist automatically classifies each query as semantic, metadata, or hybrid and executes the appropriate retrieval strategy, as described in the *Query Execution* section— investigators do not choose between them. What remains within the investigator's control is how the question itself is phrased, and phrasing has a direct effect on how precisely FTK Assist can retrieve and construct an answer. This appendix offers practical, investigation-specific guidance for phrasing queries that return sharper, more relevant results across common case types.

8.1 Anchor Queries in Case-Specific Detail

The single largest factor in result quality is how specifically a query is anchored to entities that exist in the case itself — custodian names, project codenames, account identifiers, domains, departments, and date ranges. A generic question forces the underlying retrieval to guess at scope, which tends to pull in a wider, noisier slice of the evidence than intended.

Scenario	Less Effective	More Effective
Employee misconduct	*Was anyone behaving inappropriately at work?*	*Did [Employee] make any comments about [Coworker] in chats between March and June 2026?*
Intellectual property theft	*Did someone steal our IP?*	*Did [Employee] send design specification files to a personal or external email address in the two weeks before [Resignation Date]?*

8.2 Frame Semantic Questions Around Meaning, Not Keywords

Semantic retrieval matches against the meaning of the content, so it responds best to a question phrased the way an investigator would ask a colleague who had already reviewed the evidence — naming the behavior, relationship, or intent in question, rather than a bare topic or keyword.

Scenario	Less Effective	More Effective
Insider threat	*credentials*	*Did any chat conversations suggest [Employee] was planning to access systems after [Termination Date]?*
Communications analysis	*problems with vendor*	*What concerns did the finance team raise internally about the vendor contract before it was signed?*

8.3 Frame Metadata Questions Around Countable Facts

Metadata queries are resolved as structured data requests, so they are most reliable when phrased around what can be counted, dated, or compared — message volumes, participant pairs, time windows — rather than around interpretation. As noted in Section 3.2, this retrieval scope covers emails and chat conversations.

Scenario	Less Effective	More Effective
Fraud	*Is there a lot of suspicious email activity?*	*How many emails did [Employee] exchange with [external domain] between January and March 2026?*
Insider threat	*Was [Employee] communicating a lot after hours?*	*How many chat messages did [Employee] send between 10 PM and 6 AM in the 30 days before [Resignation Date]?*

8.4 Combine Scope and Narrative in a Single Hybrid Question

Some questions genuinely need both a filter and an explanation at once — most often correspondence between two custodians. Including both the scope (who, when) and the ask (summarize, explain, list) in the same question lets FTK Assist resolve both halves in a single pass, rather than requiring a follow-up.

Scenario	Effective Hybrid Query
Timeline reconstruction	*Summarize, in chronological order, all email exchanges between [Employee A] and [Employee B] during the week [Date].*
Document review	*Across all documents shared between Finance and Procurement in Q1 2026, summarize what was discussed about the vendor discount terms.*

8.5 Refine Progressively Within a Conversation

Because follow-up questions retain the context of the conversation, the most effective pattern is to start broad enough to orient, then narrow with each follow-up rather than writing one exhaustive question up front. A typical progression during a fraud investigation:

- "What communications discuss the Q3 vendor discrepancy?"
- "Which of those involved [Employee Name]?"
- "Were there messages on this topic in the two weeks before the invoice was approved?"

Each step trades on what was already established, producing a tighter result at every stage than a single broad question could.

8.6 Reduce Noise in Results

Two habits reliably widen a query's scope beyond what was intended:

- Compound questions covering multiple threads at once, such as *"Tell me about [Employee]'s emails, chats, and any HR complaints."* Splitting this into separate questions, one topic or artifact type at a time, keeps each answer focused.
- Open-ended, whole-case questions, such as *"What's suspicious in this case?"* These draw from the entire evidence set rather than a specific line of inquiry. Naming a person, department, timeframe, or topic keeps the retrieved material relevant to the actual question being investigated.

8.7 Common Mistakes to Avoid

- Framing a question around suspicion rather than a specific act, relationship, or behavior that can be searched for in the evidence, e.g. asking "did they do anything wrong?" instead of naming what to look for.
- Using informal shorthand or nicknames unlikely to appear in the evidence itself. If a project, account, or system has a specific name in the case artifacts, use that name in the query.
- Starting a new conversation for what is really a follow-up question, losing the accumulated context that would otherwise sharpen the next answer.
- Continuing in the same conversation when the question has moved to a different custodian, topic, or investigation. Carried-over context can cause FTK Assist to anchor the new answer to entities, dates, or scope from the unrelated question that came before it — when the topic changes entirely, start fresh instead.

8.8 Query Patterns by Investigation Type

The following patterns are starting points to adapt with the specific names, dates, and terms relevant to the case at hand.

Scenario	Suggested Query Pattern
Timeline reconstruction	*List, in chronological order, all [emails/chats] involving [entity] between [date] and [date].*
Communications analysis	*How many [emails/chats] were exchanged between [Person A] and [Person B], and what topics recurred?*
Employee misconduct	*Did [Employee] make comments about [topic/person] in [artifact type] during [date range]?*
Fraud	*Summarize any communications between [Employee] and [External Party] discussing [transaction/topic].*
Insider threat	*Did [Employee]'s communications in the [period] before [departure/incident] reference [system/data/plan]?*
IP theft	*Did [Employee] share [document type] with any external or personal email address before [date]?*
Document review	*Summarize what [Department A] and [Department B] discussed about [topic] across documents in [date range].*

Contact Exterro

If you have any questions, please refer to this document, or any other related materials provided to you by Exterro. For usage questions, please check with your organization's internal application administrator. Alternatively, you may contact your Exterro Training Manager or other Exterro account contact directly.

For technical difficulties, support is available through support@exterro.com.

Contact:

Exterro, Inc.

2175 NW Raleigh St., Suite 110

Portland, OR 97210.

Telephone: 503-501-5100

Toll Free: 1-877-EXTERRO (1-877-398-3776)

Fax: 1-866-408-7310

General E-mail:info@exterro.com

Website: www.exterro.com