exterro®

EXTERRO PRODUCT LINE

EXTERRO ARMOUROP USER GUIDE

A complete walkthrough of evidence ingestion, encryption detection, the triage console, multimedia and document review, email and chat analysis, browser review, the AI assistant, bookmarks, export, and settings.

DOCUMENT TYPE	PRODUCT	PREPARED FOR	VERSION · DATE
User Guide	Exterro ARMOURop	Forensic Investigators	v1.0 · July 2026

CONTENTS

PART 12 - EXPORT

PART 13 - SETTINGS

PART 1

HOME SCREEN

Live drive triage, lab evidence ingestion, encryption handling, and job history

Exterro ARMOURop is designed for two distinct investigation scenarios: triaging evidence on a live machine at the scene, and triaging already-collected evidence in a lab. Both workflows start from the same Home screen.

Live Evidence Triage. When responding to an incident or executing a search warrant, investigators can run ARMOURop directly from a USB or external drive on the live machine. No installation on the target machine is required. The tool operates in read-only mode, meaning it never modifies the target machine or its drives. Triage can begin while the machine is still running, so there is no need to power down or image the device before getting results. The connected drives are detected automatically and shown as cards on the Home screen. Select the drive you want to triage and the process begins immediately.

Lab Triage. When evidence has already been collected and brought back to the lab, ARMOURop can be pointed directly at the collected evidence. No live machine is needed. Supported evidence types include forensic disk images (E01), mobile extraction reports (UFDR), ZIP archives, local folders, and network shares. Enter the path to your evidence in the ingestion bar and click Triage to begin. This workflow is suited for cases where evidence was collected by another team, extracted using a separate tool, or transferred from a remote location for centralised analysis.

WHY THIS MATTERS

Most triage tools require you to choose between live acquisition and lab analysis. ARMOURop handles both from a single interface, so the same workflow and the same AI-powered analysis applies regardless of where or how the evidence was collected.

The sections below walk through each part of the Home screen in detail. Start with Drive Detection if you are triaging a live machine, or Evidence Ingestion if you are working with already-collected evidence in a lab.

Drive Detection

For live triage. ARMOURop automatically detects connected drives and lets you triage directly from the card, no path entry needed.

Evidence Ingestion

For lab triage. Point the tool at a folder, network share, E01 image, UFDR, or ZIP and click Triage to begin analysis.

Encryption Detection

ARMOURop automatically detects BitLocker-encrypted volumes and encrypted iTunes backups, and prompts for credentials before triaging them.

Recent Jobs

Every triage job, past and present, is listed here with status, file count, and a View button to open results in the Triage Console.

Drive detection

ARMOURop automatically detects any drives connected to the workstation and displays them as cards below the evidence ingestion bar. You can triage directly from a detected drive without typing a path.

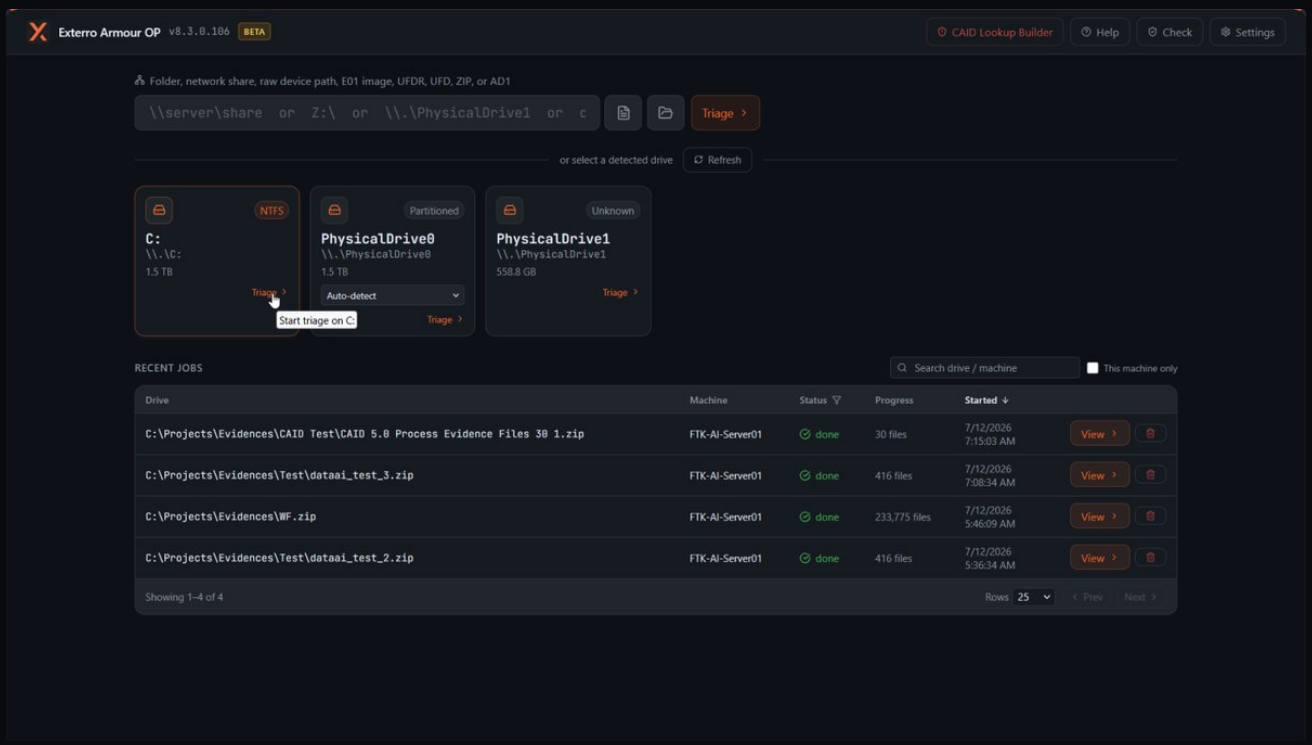

Each drive card shows the drive name, UNC path, file system type (NTFS, Partitioned, or Unknown), and total storage size. For partitioned drives, a dropdown lets you select the target partition before triaging.

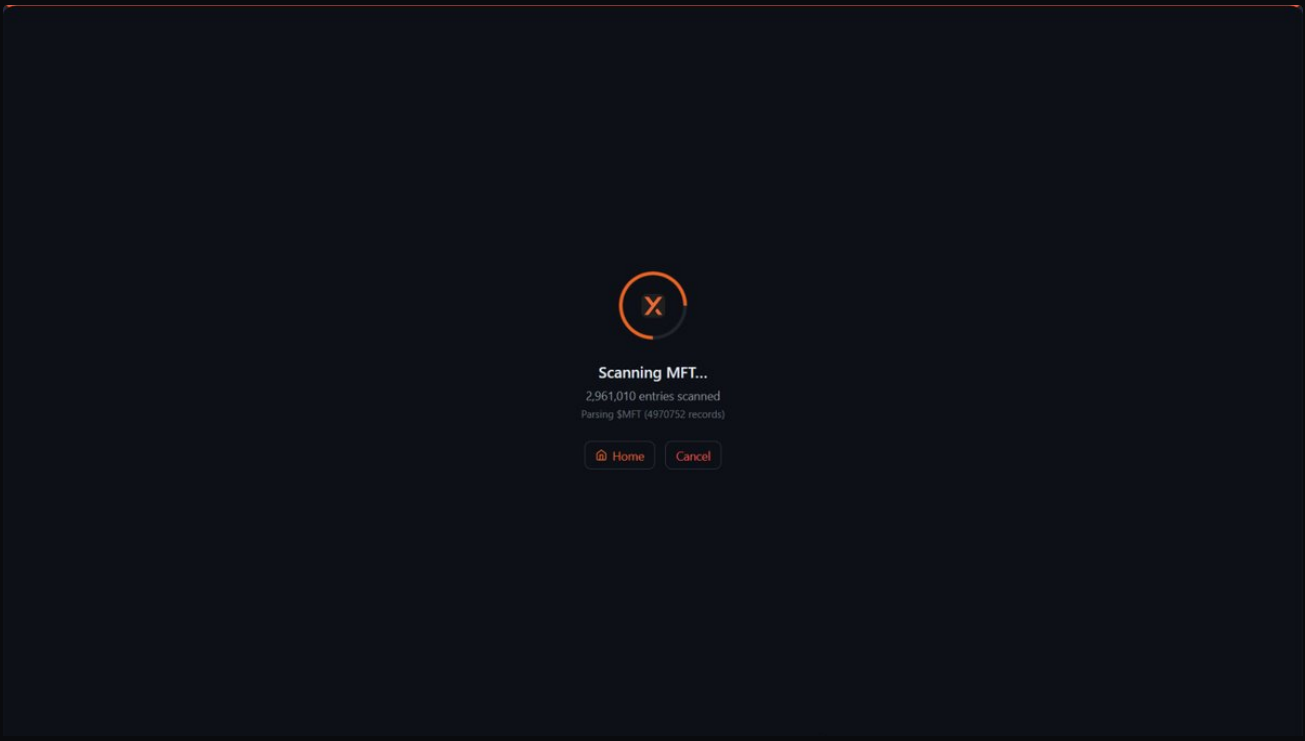

Once you click Triage on a drive card, ARMOURop begins scanning the Master File Table (MFT) to catalog all files on the drive. The progress screen shows the number of entries scanned in real time. You can click Home to return to the home screen while the job continues in the background, or Cancel to stop the job.

1. Review detected drives Connected drives appear automatically. If a newly connected drive is not showing, click Refresh to update the list.

2. Select partition if applicable For partitioned drives, use the dropdown to choose the target partition or leave it on Auto-detect.

3. Click Triage on the drive card The triage process starts immediately. No path entry needed.

> Clicking Home during triage does not cancel the job. The job continues running and will appear in the Recent Jobs table when complete.

Evidence ingestion

For lab triage, use the evidence ingestion bar at the top of the Home screen to point ARMOURop at already-collected evidence. The tool accepts a wide range of evidence formats without any conversion needed. You can point ARMOURop at evidence in two ways: Folder path A local or network folder path containing extracted evidence files.

Network share A UNC path to a network share, e.g. \server\share . Supported evidence file types for lab triage: E01 image .e01 A forensic disk image in EnCase E01 format. Split raw image .001 The first segment of a split raw disk image; subsequent segments are picked up automatically. Raw / An unformatted sector-by-sector disk image. dd .dd , .img , .raw image

Ufdr

.ufdr A mobile extraction report from Cellebrite UFED.

Ufd

.ufd An earlier Cellebrite UFED extraction report format. ZIP archive .zip A compressed archive containing evidence files. AD1 image .ad1 A logical forensic image in FTK AD1 format. TAR archive .tar An uncompressed tape archive containing evidence files. You can type or paste a path directly into the input field, or use the two icon buttons next to it to browse for evidence.

Clicking either icon first tries to open the Windows native file picker. While it loads, a pop-up appears with a fallback option in case the native picker does not show up on screen.

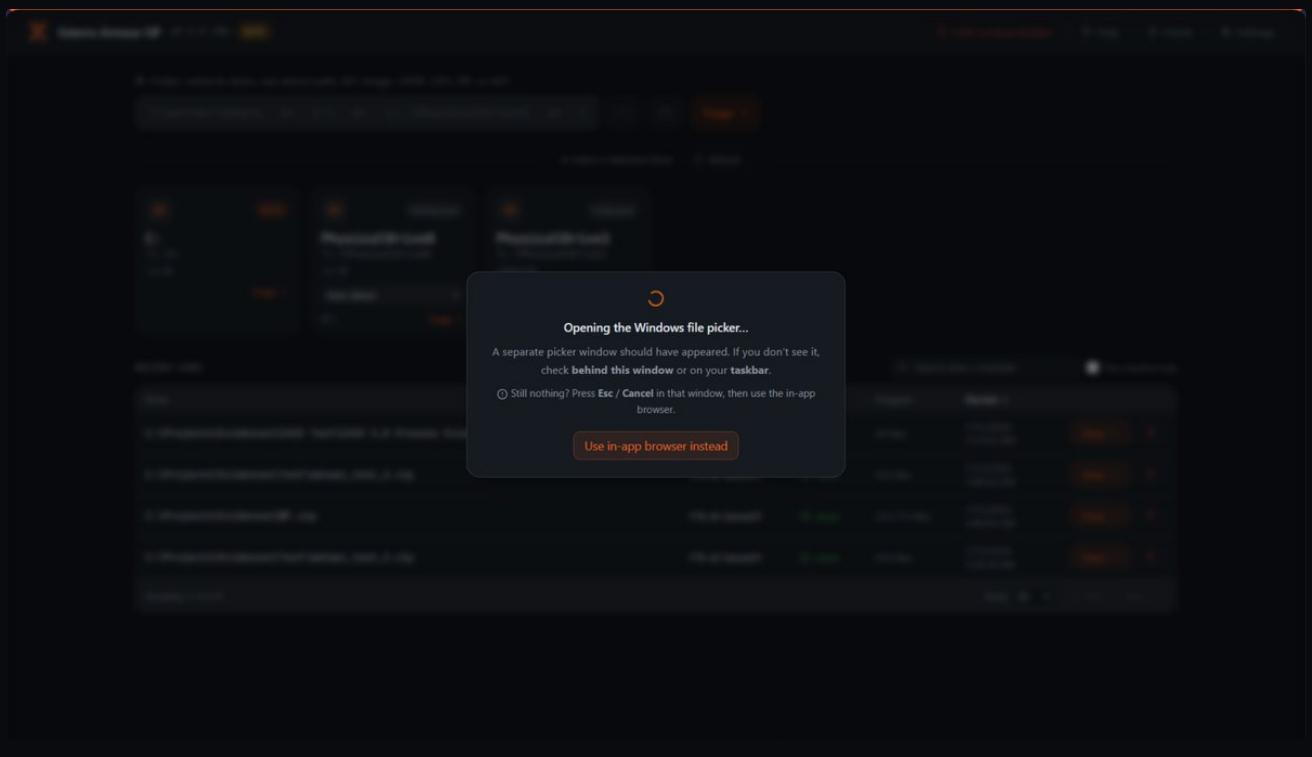

If the native picker window does not appear behind the ARMOURop window or on your taskbar within a few seconds, press Esc or Cancel in that window, then click Use in-app browser instead to browse using ARMOURop's built-in file browser.

Browsing For A File

Use the file icon to browse for a single evidence file, such as an E01 image, UFDR, or ZIP archive.

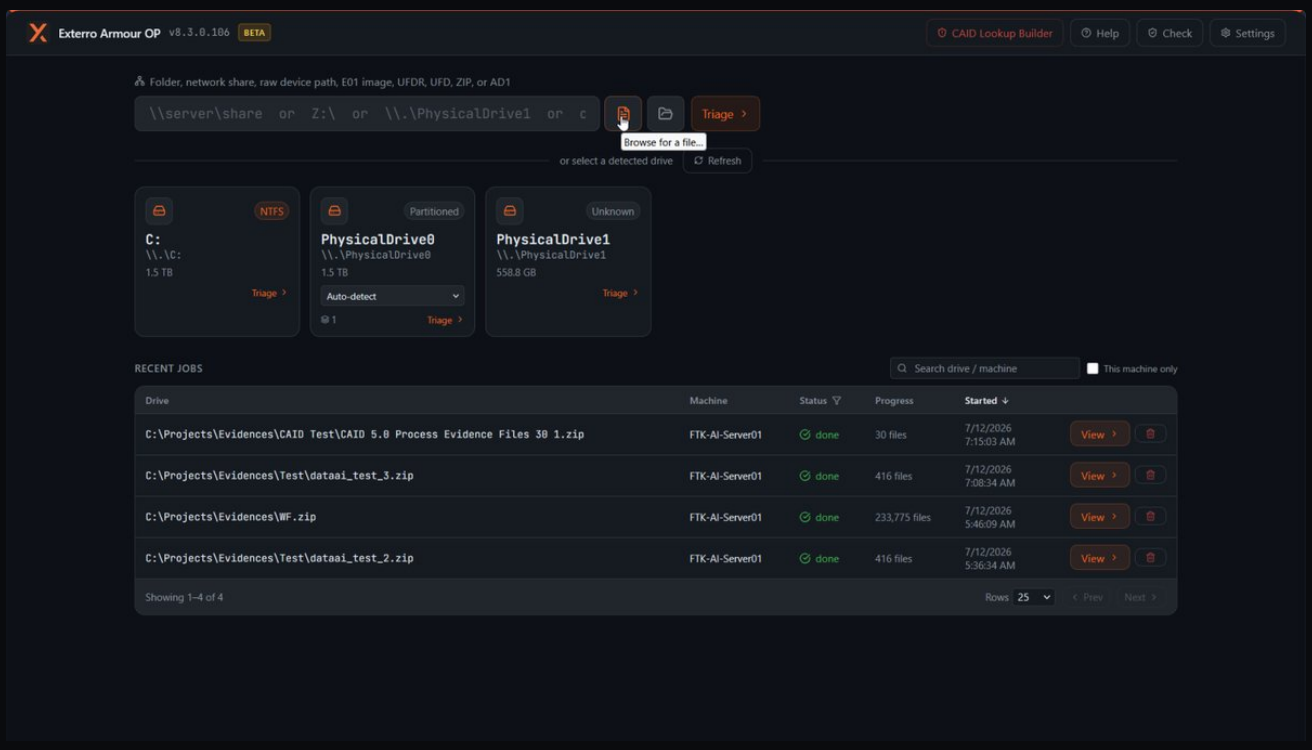

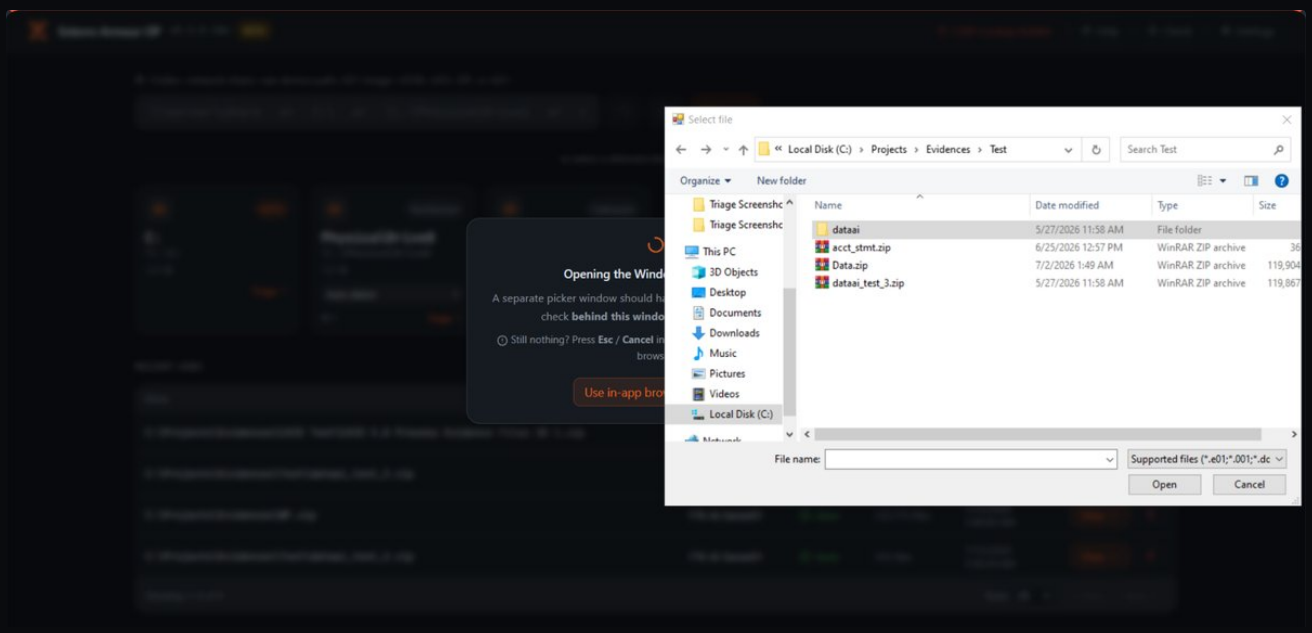

1. Click the file icon Located next to the path input field. This opens the Windows native file picker for selecting a single evidence file.

2. Select your evidence file Navigate to the file and click Open. The selected path is loaded into the input field. BROWSING FOR A FOLDER Use the folder icon to browse for a folder of extracted evidence files or a network share.

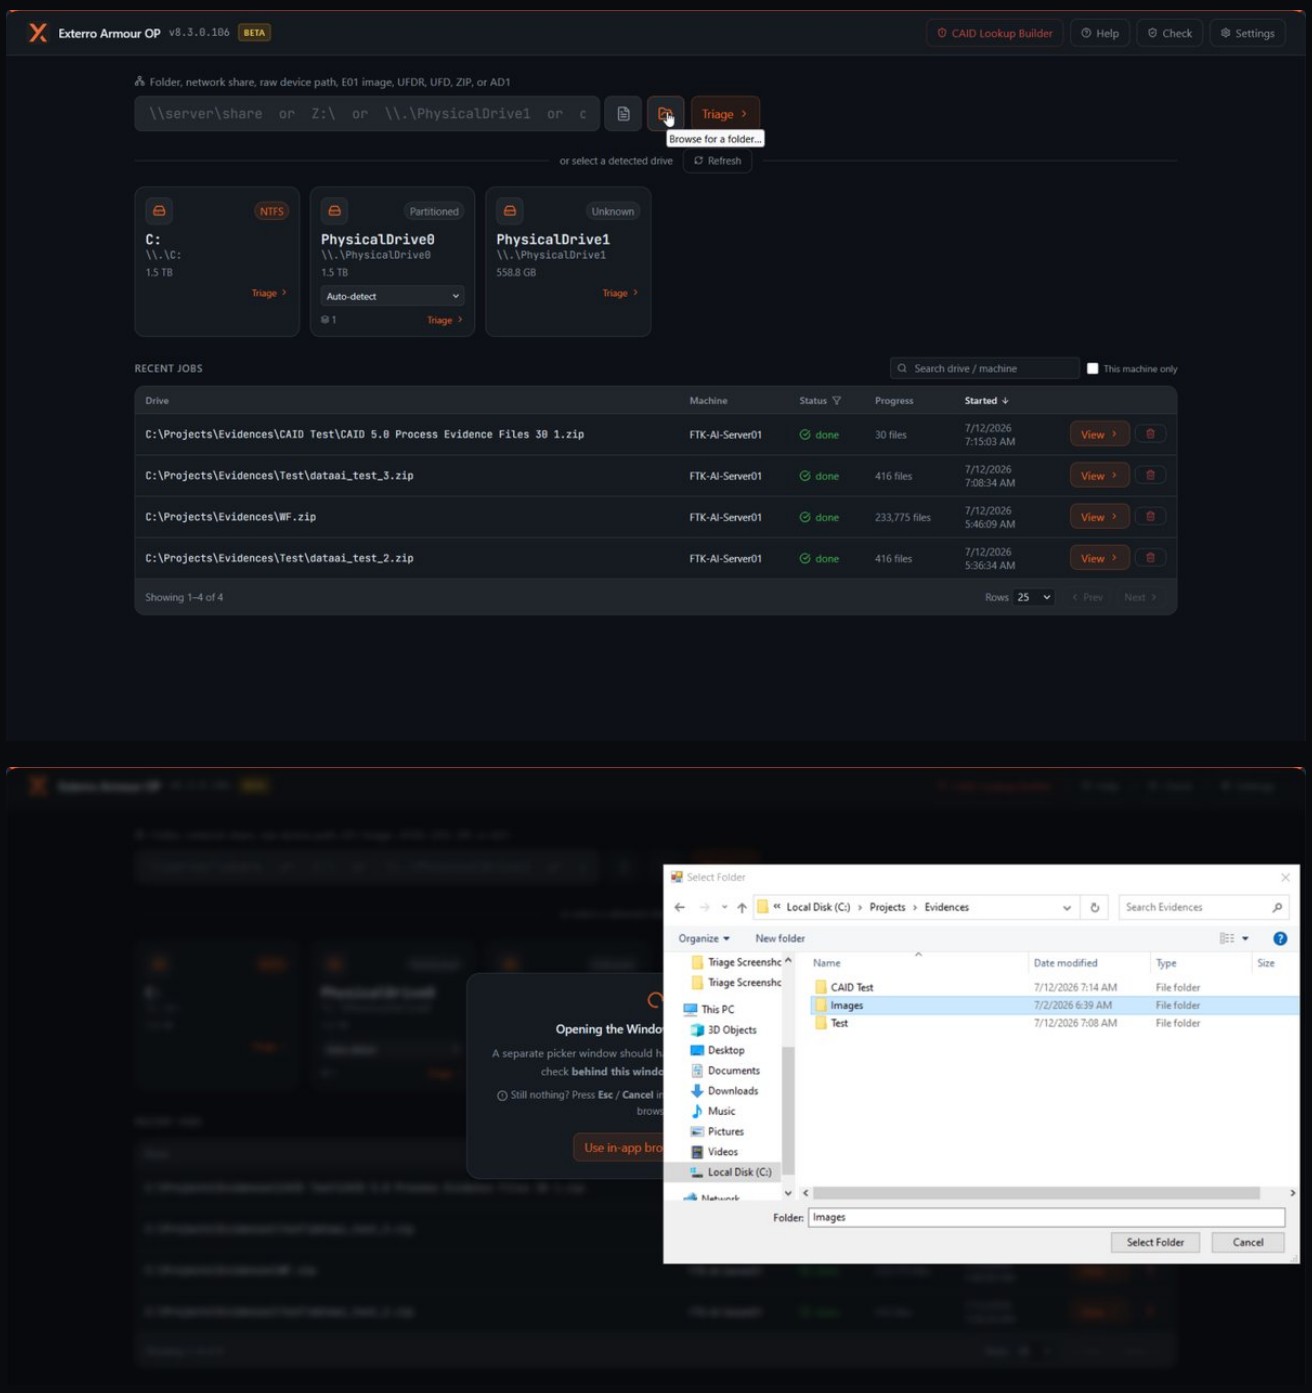

1. Click the folder icon Located next to the path input field. This opens the Windows native folder picker.

2. Select your evidence folder Navigate to the folder and click Select Folder. The selected path is loaded into the input field. USING THE IN-APP BROWSER INSTEAD If the Windows native picker does not appear, click Use in-app browser instead from the loading pop-up. This opens ARMOURop's own file browser, which works the same way for both files and folders.

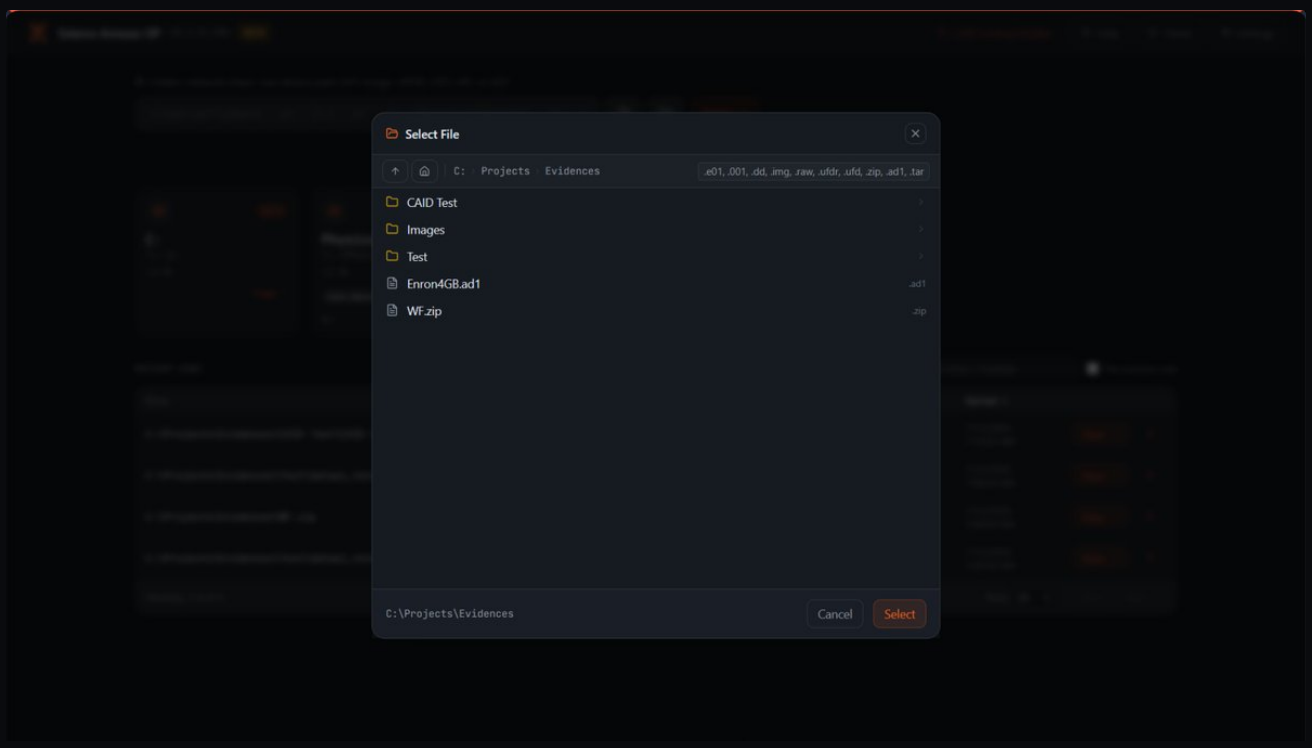

The in-app browser lists only file types supported for triage, shown in the extension filter bar at the top. The full selected path is shown at the bottom of the window before you click Select.

1. Click Triage Once your evidence path is loaded into the input field, click the orange Triage button to start the AI-powered triage process. The job appears in the Recent Jobs list with live status updates.

Multiple triage jobs can run simultaneously. Each job is tracked independently in the Recent Jobs list, so you can monitor several cases at the same time.

Encryption detection

ARMOURop automatically detects encrypted evidence during triage and pauses to ask for credentials before continuing. Two encryption types are currently detected: BitLocker encrypted volumes and encrypted iTunes backups.

Bitlocker Encryption Detection

When a disk image or drive contains one or more BitLocker-encrypted volumes, ARMOURop pauses triage and opens the BitLocker Encryption Credentials dialog. For each encrypted volume, you can provide a boot key file, recovery key, or user password.

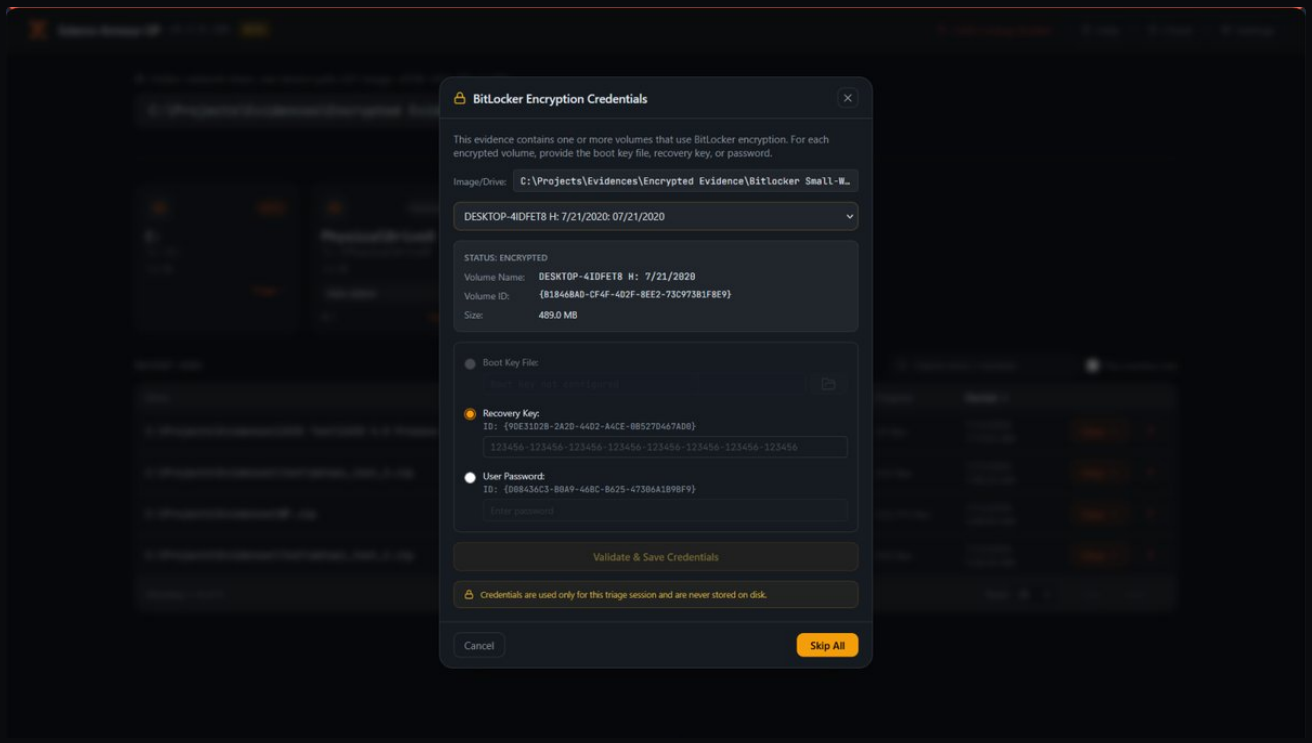

If the image or drive contains multiple BitLocker-encrypted volumes, a dropdown lets you switch between them and provide credentials for each one individually before continuing.

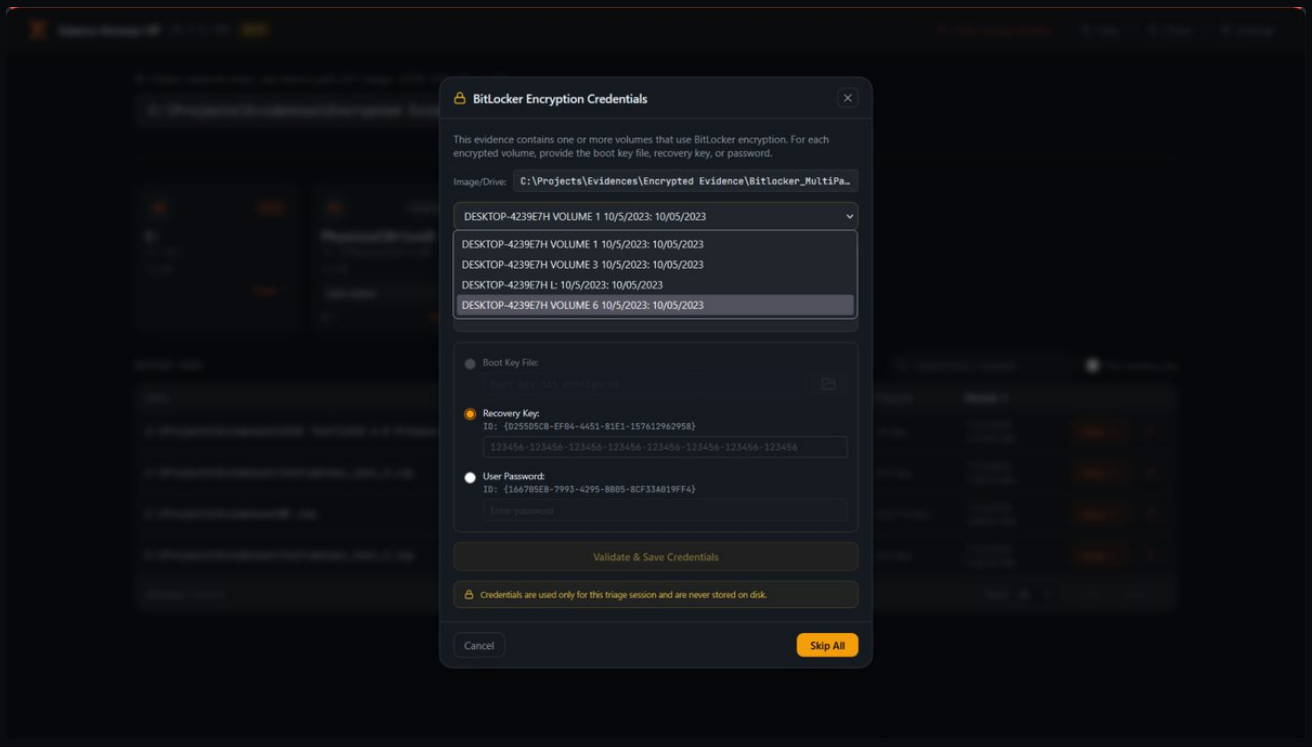

1. Review the volume details The dialog shows the Image/Drive path, encryption status, volume name, volume ID, and size for the volume currently selected.

2. Choose a credential type Select Boot Key File and browse to a .bek file, enter the Recovery Key (48-digit numerical key), or enter the volume's User Password.

3. Validate & Save Credentials If there is more than one encrypted volume, repeat for each volume using the dropdown, then click Validate & Save Credentials to unlock the volumes and continue triage.

> Cancel stops the triage job entirely. Skip All continues triage without decrypting any of the encrypted volumes on this image or drive - encrypted volumes are skipped rather than indexed. Credentials you enter are used only for this triage session and are never stored on disk.

Encrypted Itunes Backup

When ARMOURop encounters a password-protected iTunes backup within the evidence, it pauses and opens the Encrypted iTunes Backup dialog, prompting for the backup password.

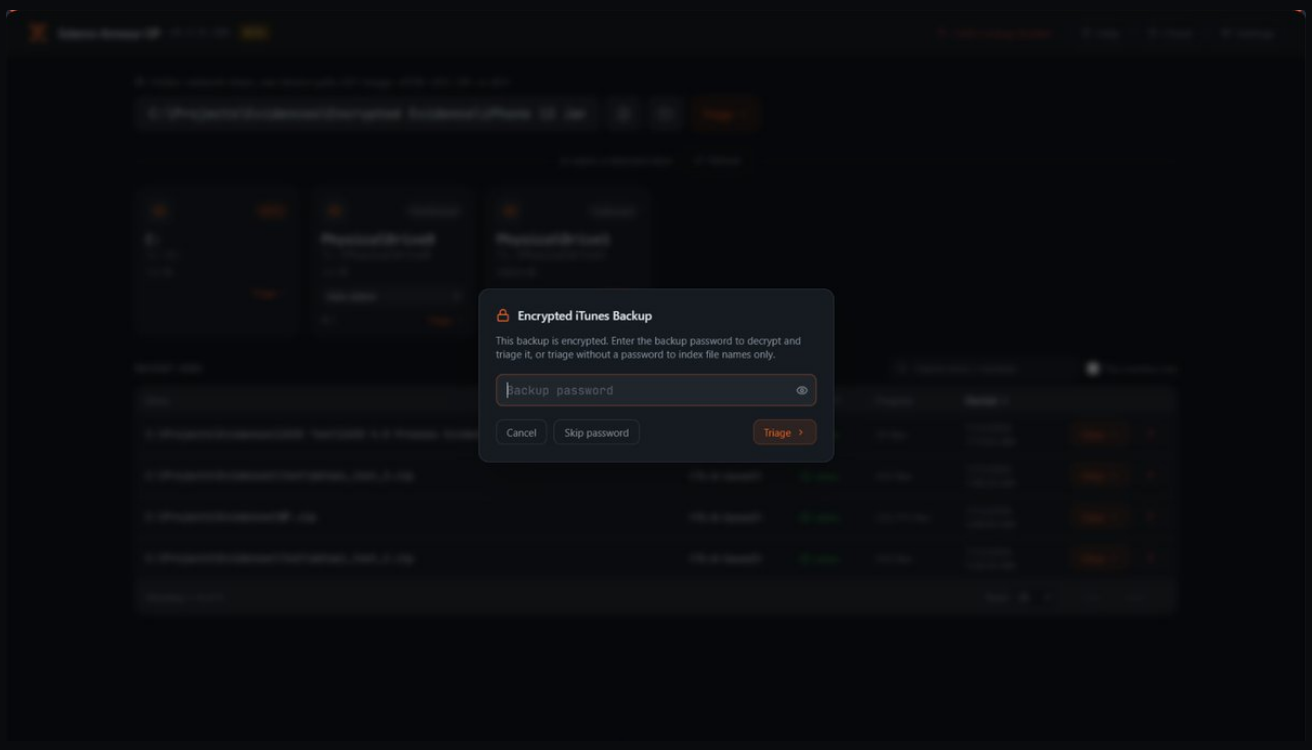

1. Enter the backup password Type the password used to encrypt the iTunes backup into the Backup password field. Click the eye icon to reveal the typed password.

2. Click Triage Decrypts the backup with the provided password and continues triage with full access to its contents.

> Cancel stops the triage job entirely. Skip password continues triage without decrypting the backup - only file names are indexed, file contents remain inaccessible.

Recent jobs

Every triage job, past and present, appears in the Recent Jobs table at the bottom of the screen. This gives you a full audit trail of what has been triaged, when, and how many files were processed.

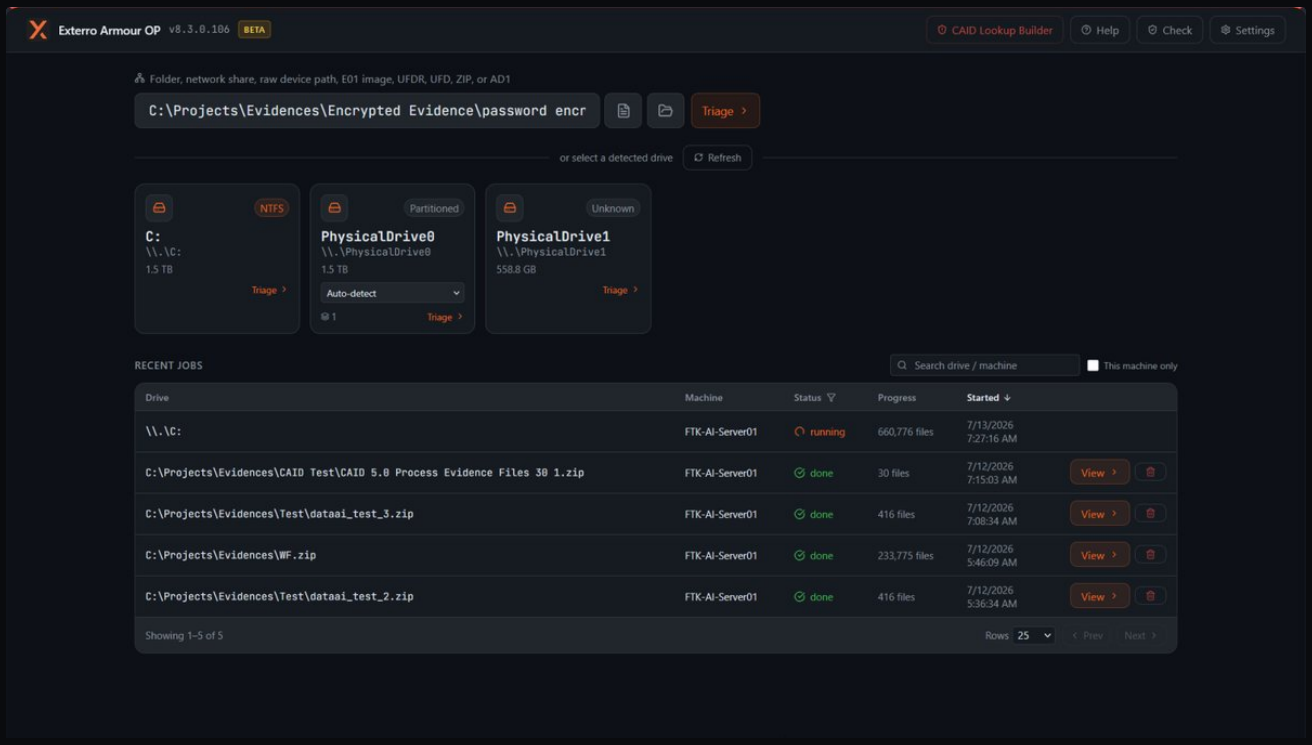

Each row in the table shows: Drive The full path to the evidence source that was triaged. Machine The name of the machine the triage job ran on. Status Running for active jobs, Done for completed jobs. Click the column header to filter by status. Progress The number of files processed so far, or the total file count for completed jobs.

Started The date and time the triage job began. Click the column header to sort by start time. Use the Search drive / machine box above the table to filter jobs by evidence path or machine name. Check This machine only to show jobs run on the current machine only.

1. Monitor a running job Jobs with a Running status show a live file count as triage progresses. You can navigate away and return at any time without interrupting the job.

2. View completed results Click the View button on any completed job to open the Triage Console and explore all discovered artifacts.

3. Delete a job Click the trash icon to remove a job record from the list. This does not delete the original evidence.

> Jobs are listed with the most recently started at the top. A job showing Running can be left to complete in the background while you review results from a previous job.

PART 2

GPU SETUP

Verifying and staging CUDA / cuDNN for GPU-accelerated AI features

ARMOURop can use an NVIDIA CUDA-enabled GPU to accelerate its AI-powered features. This guide covers the two NVIDIA components you need - CUDA 12 and cuDNN 9 - and how to confirm ARMOURop can see them. GPU acceleration is optional. ARMOURop runs fine without a GPU - GPU-accelerated features are simply disabled.

Follow this guide to enable them on machines with a supported NVIDIA GPU.

What You Will Need

A machine with an NVIDIA CUDA-capable GPU and administrator rights to install the NVIDIA components. STEP 1 - INSTALL THE NVIDIA CUDA TOOLKIT 12 ARMOURop's GPU build targets the CUDA 12 runtime. Install the latest 12.x release (for example, 12.9.x).

> Pick CUDA 12, not 13. NVIDIA's main download page now defaults to CUDA 13. ARMOURop requires CUDA 12.x - use the CUDA Toolkit Archive and select a 12.x release for Windows.

1. Open the CUDA Toolkit Archive Navigate to the CUDA Toolkit Archive on developer.nvidia.com and select the latest 12.x release for Windows.

2. Run the installer Accept the Express (default) options. A reboot may be requested once installation completes. STEP 2 - INSTALL CUDNN 9 cuDNN 9 provides the deep-learning libraries ARMOURop uses on the GPU. Install the latest 9.x release built for CUDA 12.

> An NVIDIA developer account (free) may be required to download cuDNN. When prompted during installation, choose the cuDNN 9 package for CUDA 12.x, not CUDA 13.

1. Open the cuDNN Downloads page Go to the cuDNN Downloads page on developer.nvidia.com and select the Windows build for CUDA 12.x.

2. Run the cuDNN installer Accept the defaults and complete the installation. STEP 3 - VERIFY THE GPU IN ARMOUROP With CUDA 12 and cuDNN 9 installed, confirm that ARMOURop can see your GPU.

1. Launch ForensicTriage.exe as administrator Right-click ForensicTriage.exe and select Run as administrator.

2. Click the Check button Click the Check button in the top-right corner of the home page. ARMOURop inspects your GPU, copies any drivers it needs, and opens a System Check window. The Check button, top-right

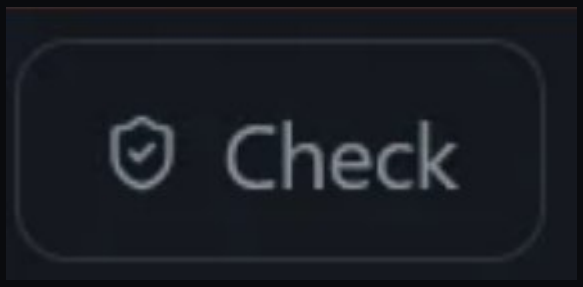

If everything is in place, the System Check window will show green ticks against NVIDIA GPU detected, CUDA / cuDNN DLLs present, and Running GPU build, plus All default models present. No further action is needed - ARMOURop is GPU-accelerated.

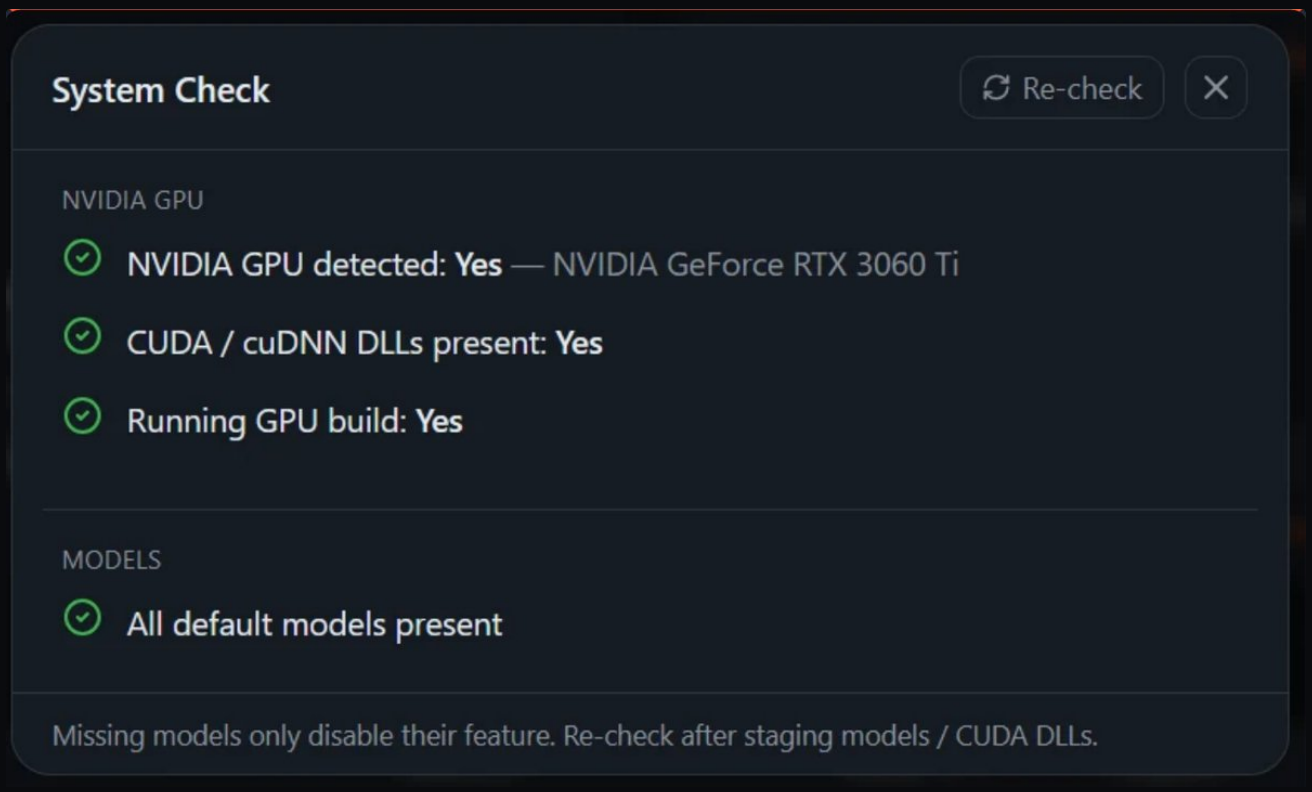

STEP 4 - IF THE CHECK REPORTS MISSING DLLS If the System Check shows a red cross against CUDA / cuDNN DLLs present: No, ARMOURop cannot yet see the cuDNN libraries. Stage them manually so the application has direct access.

1. Open the cuDNN bin folder In File Explorer, navigate to: C:\Program Files\NVIDIA\CUDNN\v9.x\bin\12.x\x64 . Substitute your installed versions - v9.x is your cuDNN version and the bin sub-folder (e.g. 12.x) is named after your CUDA version.

2. Copy all .dll files Select and copy all .dll files in that folder.

3. Create the cuda folder in ARMOURop In your ARMOURop installation folder, open forensic_triage_internal\tools\ and create a new folder named cuda if it doesn't already exist.

4. Paste the DLLs Paste the copied DLLs into the cuda folder: ...\forensic_triage_internal\tools\cuda .

5. Re-check in ARMOURop Back in the System Check window, click Re-check. The check should now report CUDA / cuDNN DLLs present: Yes and GPU-accelerated features will be ready to use.

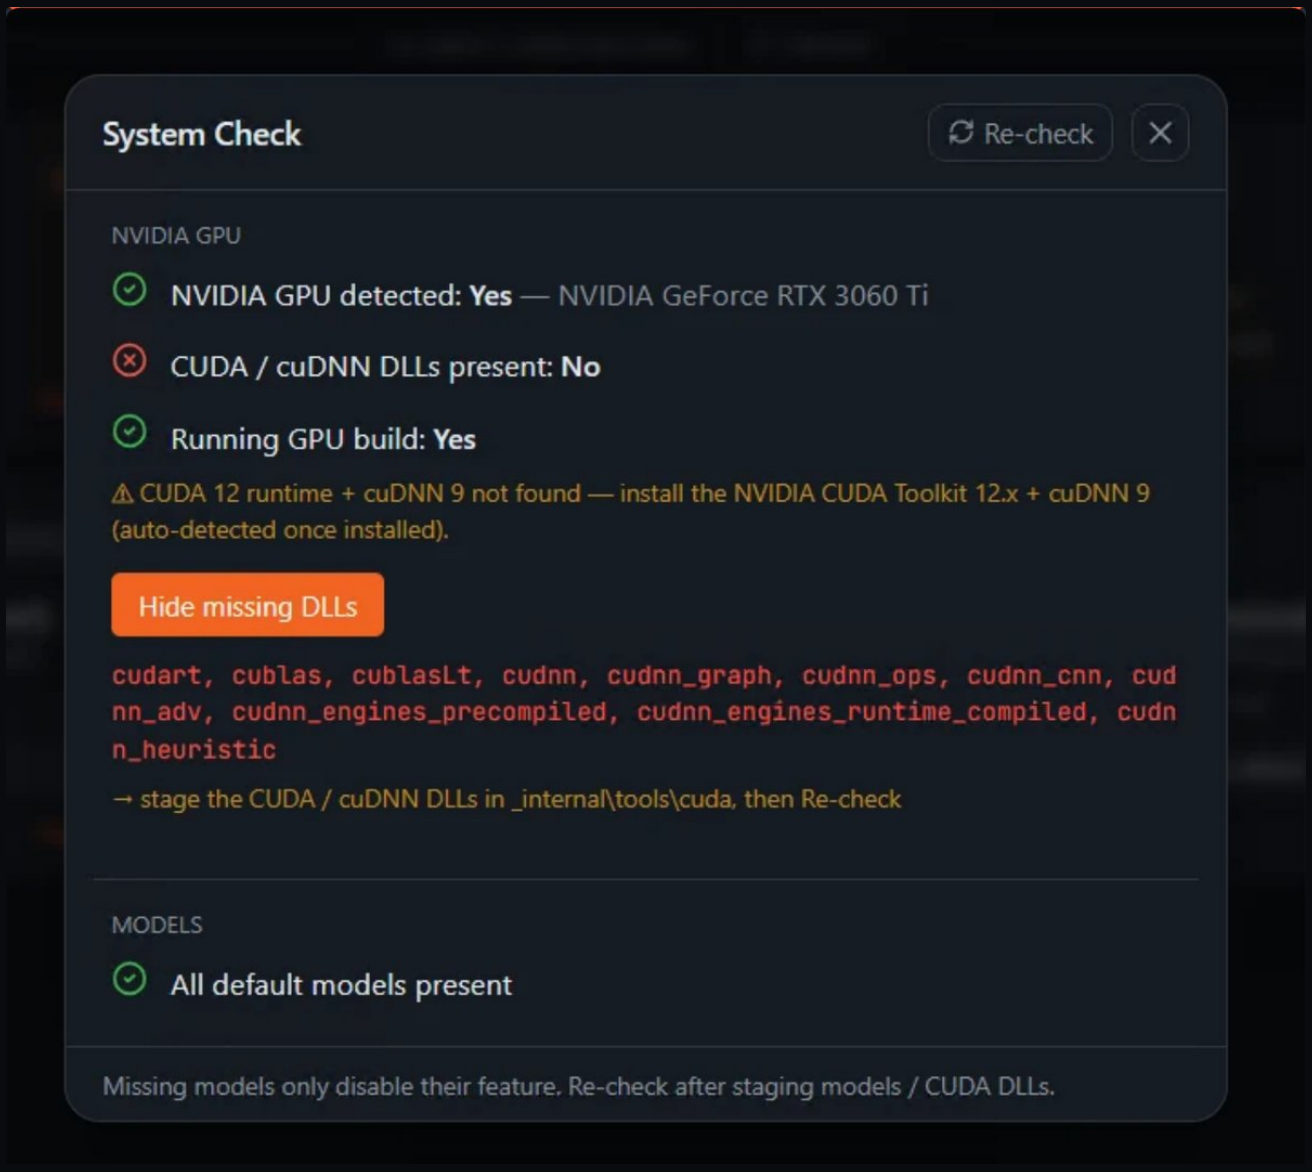

Missing models only disable their individual feature - they don't affect GPU setup. You can re-run the check at any time using the Re-check button after staging models or DLLs.

PART 3

CAID LOOKUP BUILDER

Loading known illicit image hash sets for automatic CSAM checking

The CAID Lookup Builder lets you load known illicit image hash sets into a local database. Once loaded, every triage run automatically checks discovered images against these hashes and flags any matches. Why this matters: Manually grading images for CSAM can take a week or more.

With a hash database loaded, ARMOURop checks tens of thousands of images against known hashes in hours, not days.

How To Access

From the Home screen, click the CAID Lookup Builder button in the top-right corner.

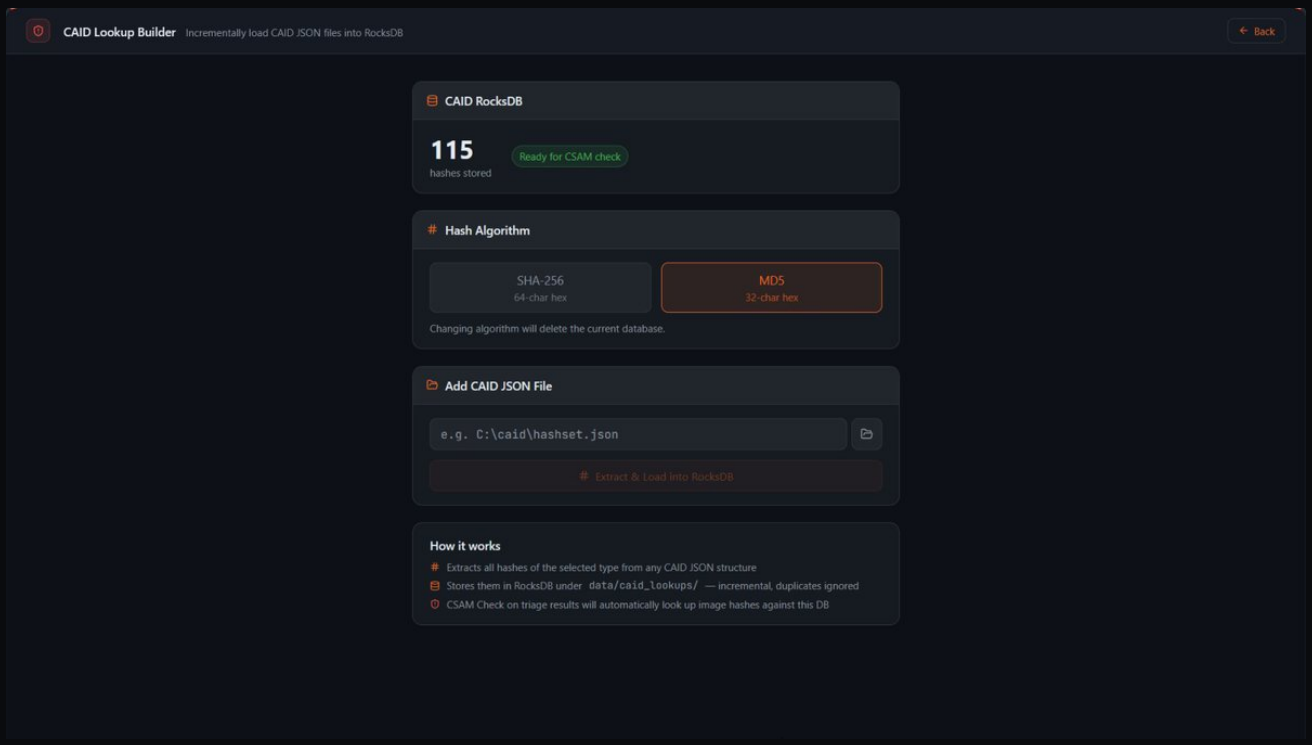

Database Status

The top panel shows how many hashes are stored in the local RocksDB database. A green "Ready for CSAM check" badge confirms the database is active and will be used automatically during triage. If setting up for the first time, the count will show zero.

Hash Algorithm

Select the algorithm that matches the format of your CAID JSON file before loading any hashes.

Sha-256

64-character hex. Select this if your CAID JSON uses SHA-256 hashes.

Md5

32-character hex. Select this if your CAID JSON uses MD5 hashes.

> Important: Switching algorithm will delete the existing database. Confirm the correct format before changing this setting.

Loading Hash Sets

1. Enter the JSON file path Type the full path to your CAID JSON file (e.g. C:\caid\hashset.json), or click the browse icon to locate it.

2. Click Extract and Load into RocksDB The tool extracts all hashes of the selected type and stores them in the local database. The hash count updates once loading is complete.

> You can add multiple JSON files over time. Each time you receive an updated hash set, repeat the loading steps. New hashes are added on top of what is already stored. Duplicates are ignored automatically.

PART 4

TRIAGE CONSOLE

Navigating file categories, folder filters, AI priority, and advanced filters

Once a triage job is complete, clicking View on the home screen opens the Triage Console. This is your primary workspace for reviewing evidence, narrowing down suspects, identifying key files, and preparing materials for the next stage of the investigation.

Why this matters: Instead of manually sifting through thousands of files, the Triage Console gives investigators a structured, filtered view of all discovered artifacts. AI-generated attributes, folder signals, and priority scores mean you can cut straight to what matters most.

How To Access

From the Home screen, find the completed job in the Recent Jobs table and click the View button.

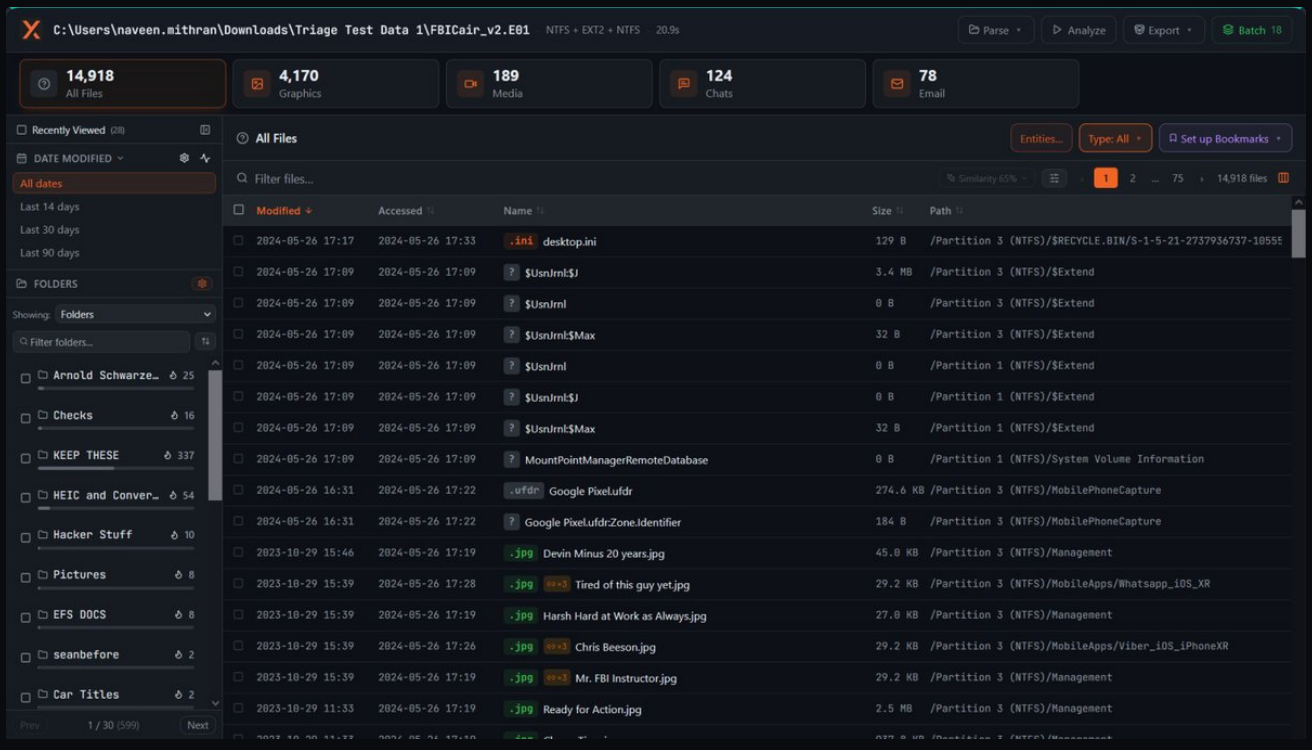

File category summary bar

The five cards across the top of the Triage Console give you an instant count of files by category. Clicking any card switches the file list to show only that category.

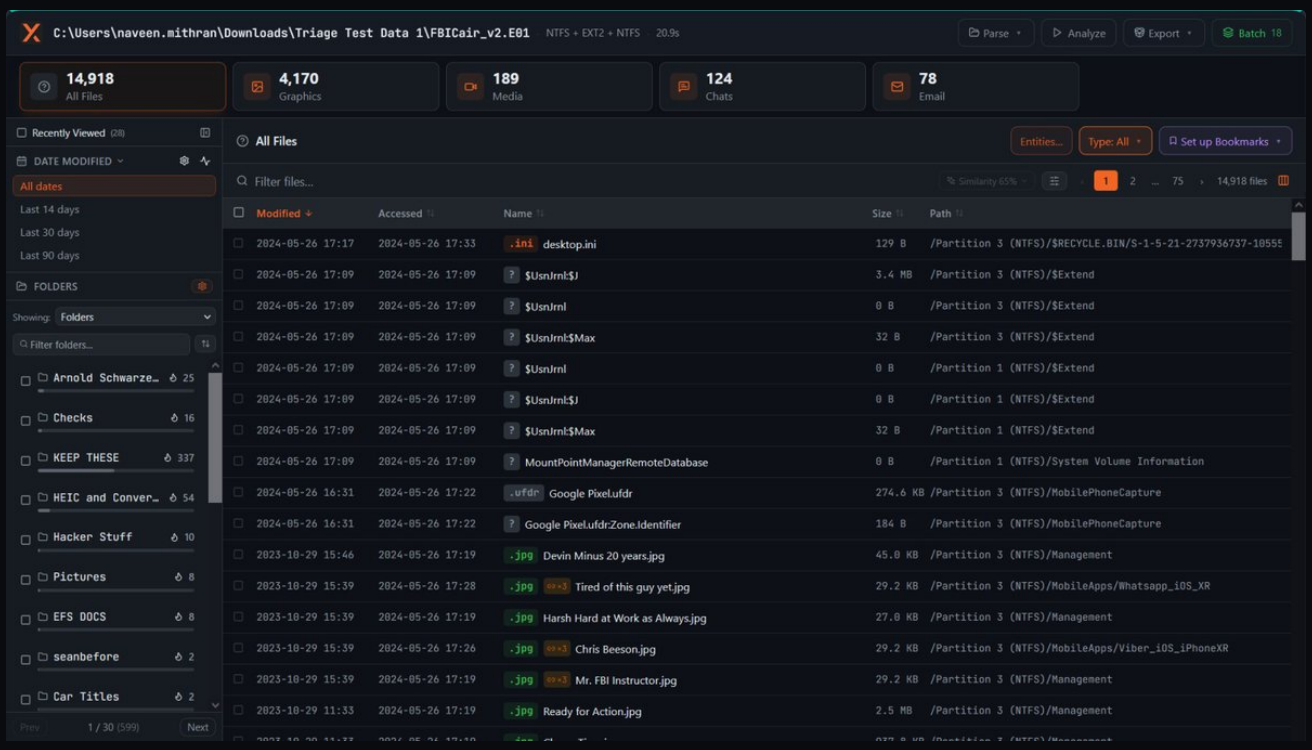

All Files Every file discovered across the entire evidence source, regardless of type. Graphics All image files. Includes AI captions, face detection results, and CSAM hash matches. Media All video and audio files. Includes transcripts, video frame previews, and in-video search. Chats Extracted chat and messaging data from the device or extraction file. Email All email data.

Searchable and filterable by folder, sender, domain, and attachments.

> Each category has its own dedicated section in this guide with detailed instructions for review and analysis.

Left panel filters

The left panel is the primary way to narrow down what you see in the file list. It has three sections: date filters, folder navigation, and AI-driven folder intelligence. Filter by date modified

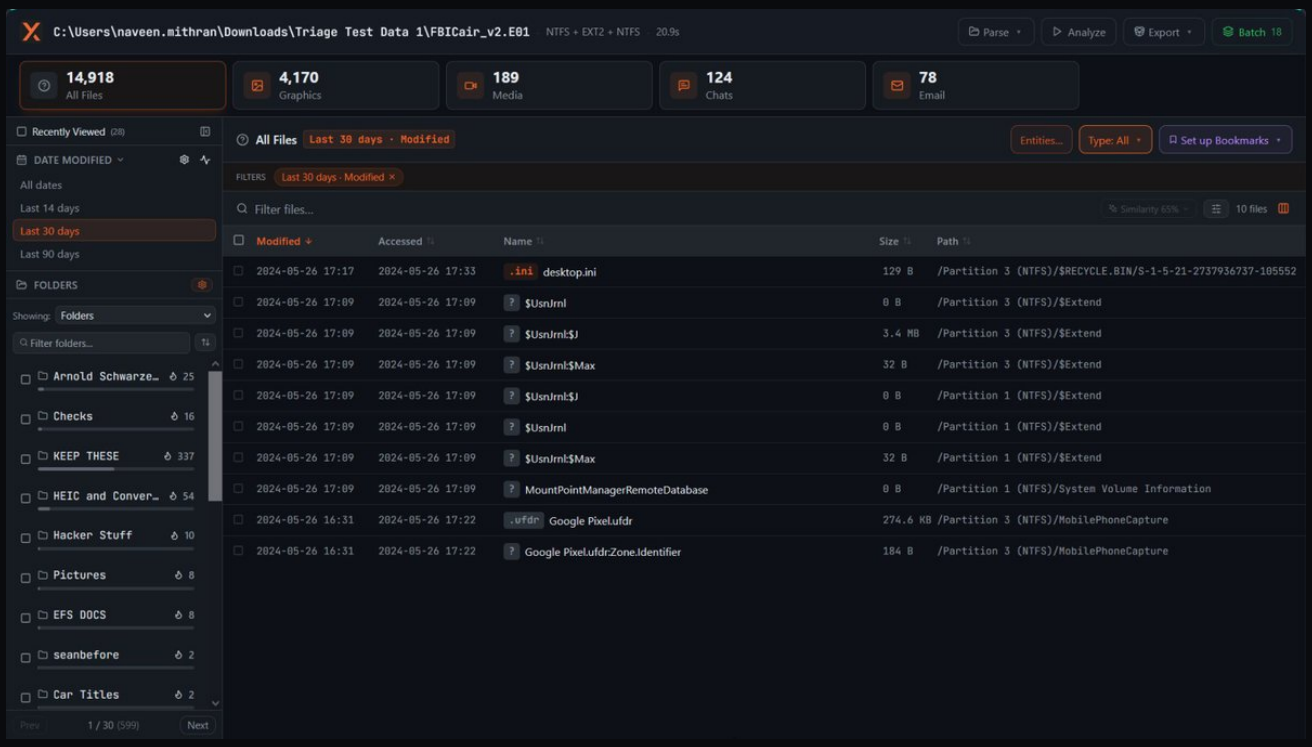

Use the Date Modified panel to quickly narrow the file list to a specific time window. Choose from All dates, Last 14 days, Last 30 days, or Last 90 days. Filter by folders

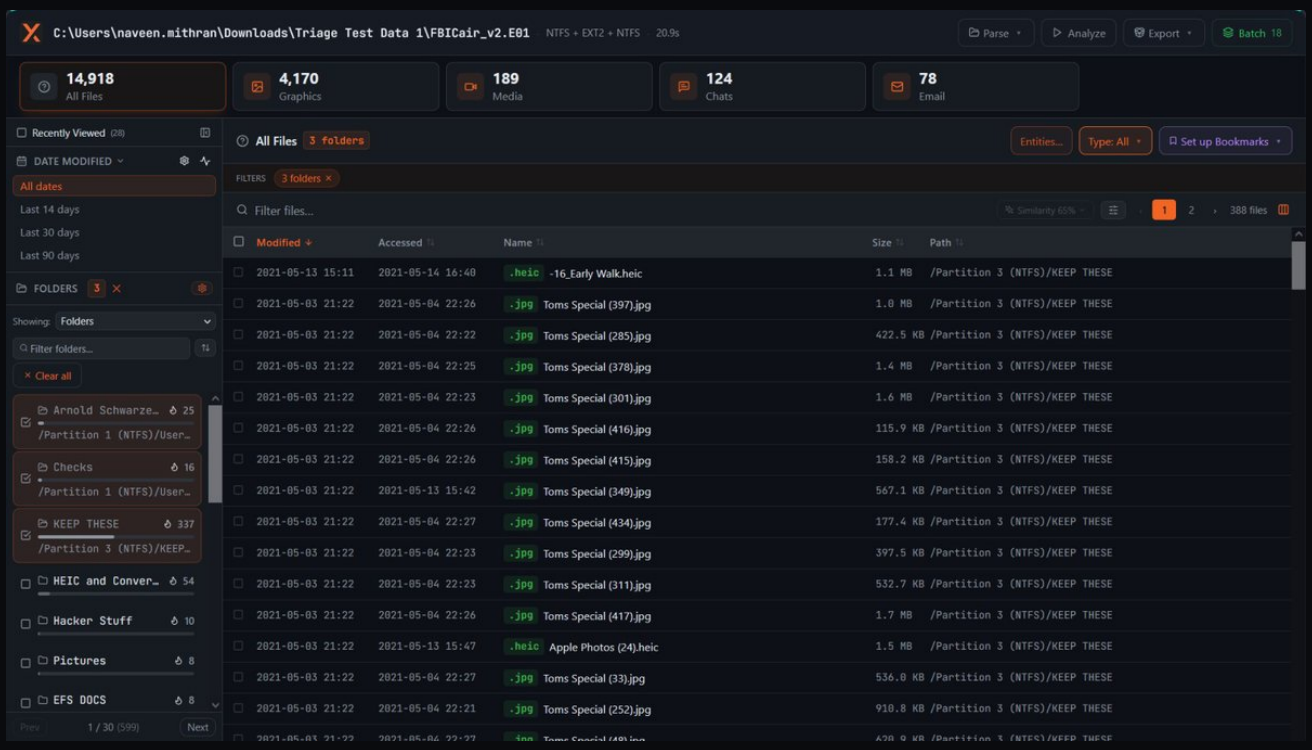

The Folders panel lists every folder found in the evidence. Each folder shows a name, file count badge, and partition path. Select one or multiple folders to focus the file list on specific locations. Sort folders

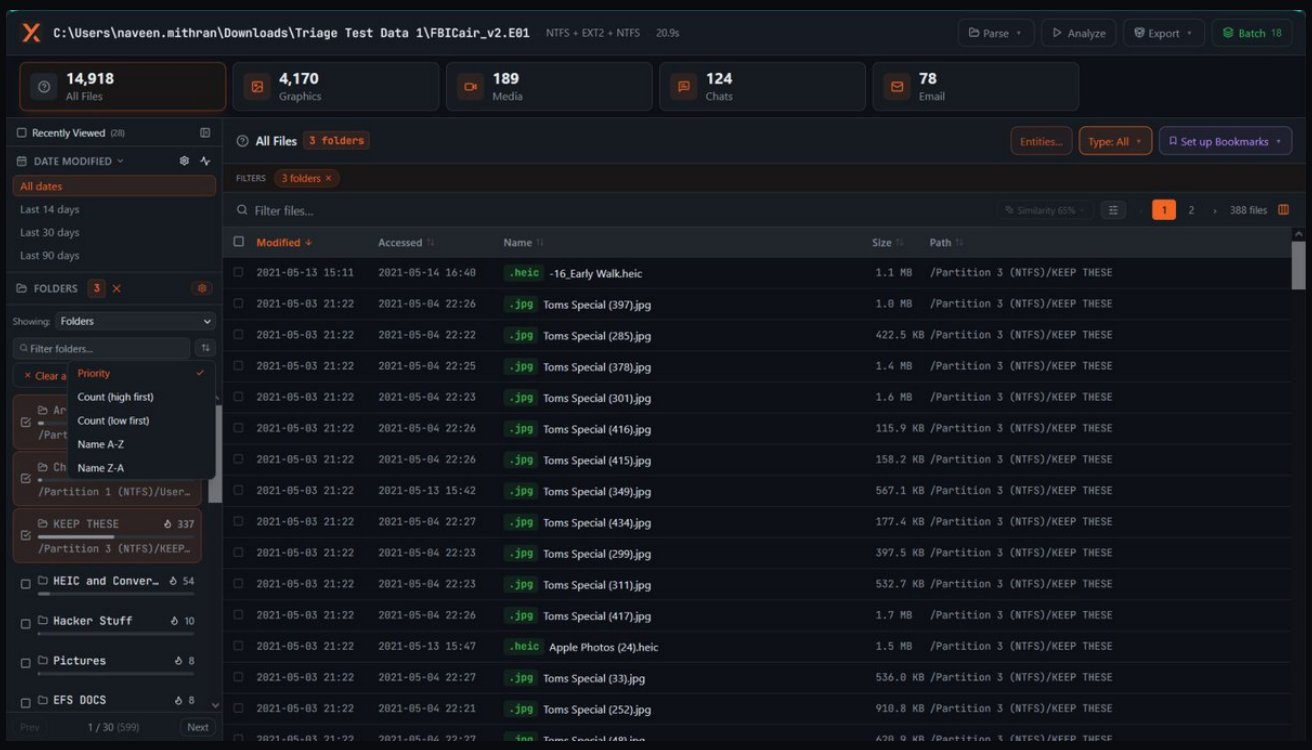

Priority Default Folders ordered by AI-assigned relevance score, most significant first. Count (high first) Folders with the most files appear at the top. Count (low first) Folders with the fewest files appear at the top. Name A-Z Folders sorted alphabetically from A to Z. Name Z-A Folders sorted alphabetically from Z to A.

AI Priority and Signals

During triage, the AI automatically evaluates every folder and assigns two types of intelligence: a Priority score and a set of Signals. These help investigators cut directly to the folders most likely to contain relevant evidence.

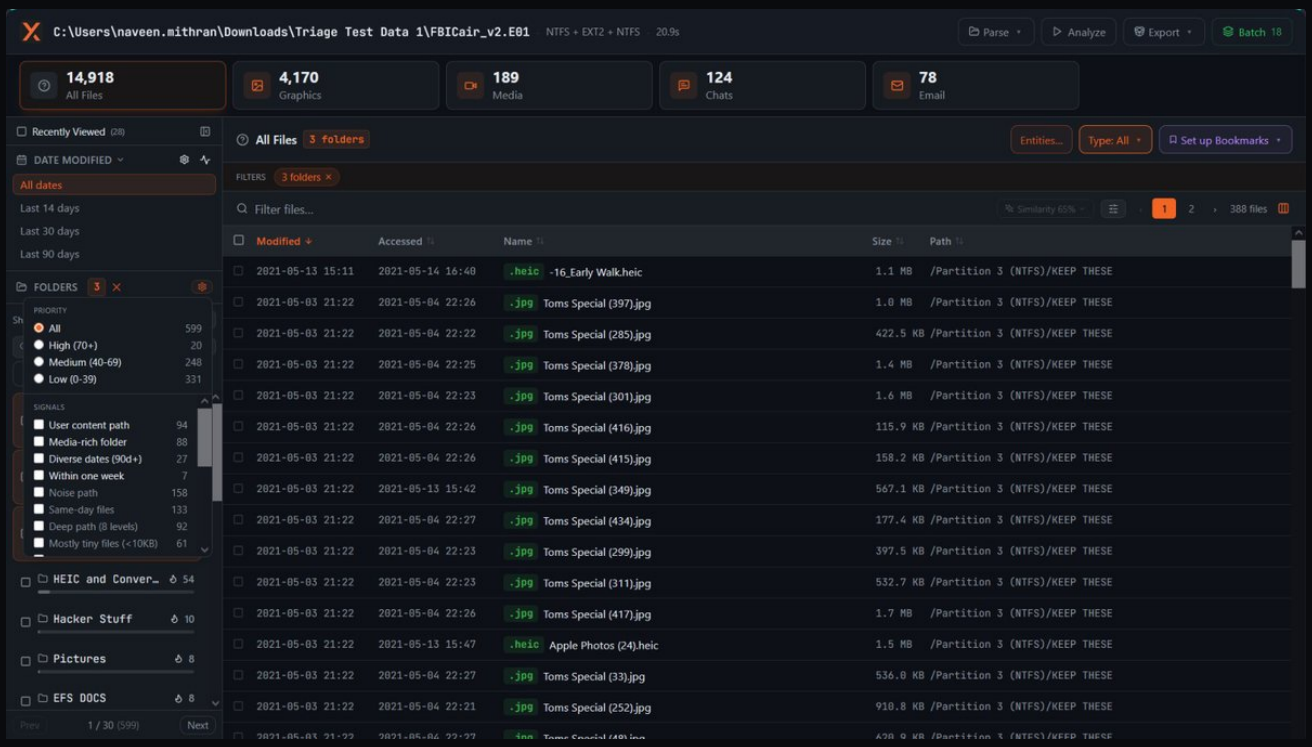

Priority scores The AI scores each folder on a scale of 0 to 100 based on investigative relevance. Use the Priority filter to focus on the folders that matter most. High 70+ Most likely to contain key evidence. Focus here first. Medium 40-69 Worth reviewing if high-priority folders are exhausted. Low 0-39 Least likely to contain relevant material.

Signals Signals are behavioral characteristics the AI detects in each folder. Use them to surface folders that match specific patterns relevant to your investigation. User content path Folder path associated with user-generated content. Media-rich folder Folder contains a high proportion of image and video files. Diverse dates (90d+) Files span more than 90 days of activity.

Within one week All files were created or modified within the same week. Noise path Folder appears to contain system or non-user files. Same-day files Most files share the same creation date. Deep path (8 levels) Folder is buried 8 or more levels deep in the directory tree. Mostly tiny files Folder is dominated by files under 10KB.

Advanced filters

The advanced filter panel gives you precise control over what appears in the file list. Filters are grouped into five categories and can be combined freely.

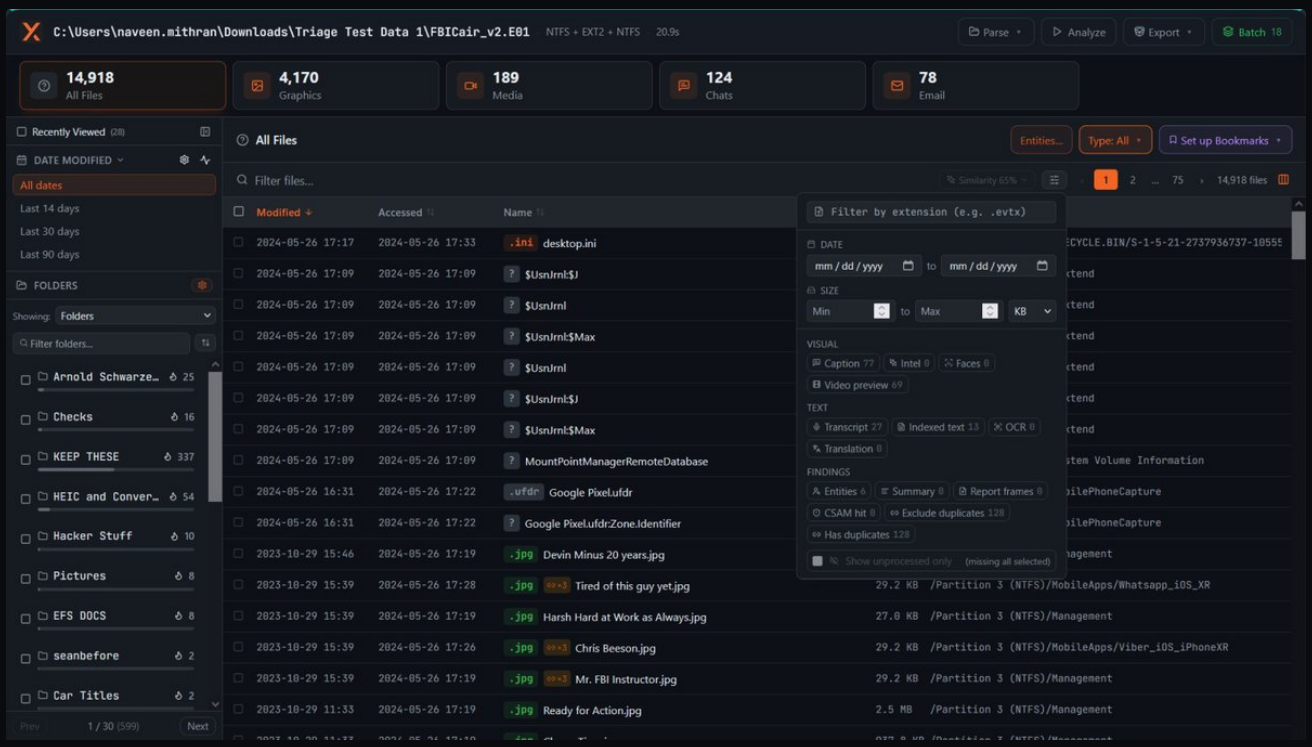

Extension Type a file extension (e.g. .evtx, .pdf, .jpg) to show only files of that type. Date Set a from and to date range to narrow the file list to a specific time window. Size Set a minimum and maximum file size to filter out files that are too small or too large. Visual Filter by AI-generated visual attributes: Caption, Intel, Faces, and Video preview.

Generated when you run the Analyze function. Text Filter by AI-generated text attributes: Transcript, Indexed text, OCR, and Translation. Generated when you run the Analyze function. Findings Filter by investigation findings: Entities, Summary, Report frames, CSAM hit, Has/Exclude duplicates, and Show unprocessed only.

> Filters from the left panel and the advanced filter panel work together. For example, select a high-priority folder, filter by Last 30 days, then narrow to only images with detected faces, all at the same time.

PART 5

MULTIMEDIA REVIEW

Reviewing images and video with AI captions, OCR, transcription, and face detection

The Multimedia Review section covers how to view, filter, and analyze images and videos within a triage job. Using built-in views and AI-powered analysis tools, investigators can surface key evidence from thousands of files in a fraction of the time it would take manually.

Why this matters: Reviewing multimedia evidence manually is one of the most time-consuming tasks in any investigation. ARMOURop uses AI to classify images, generate captions, read text from photos, transcribe audio, detect faces, and flag known illicit content automatically. WHEN WOULD I USE THIS? Any time a triage job contains images or videos.

Navigate to the Graphics or Media tab in the Triage Console to begin.

Graphics

The Graphics tab shows all images discovered in the evidence as a thumbnail grid. The same date, folder, and advanced filters from the All Files view are available here.

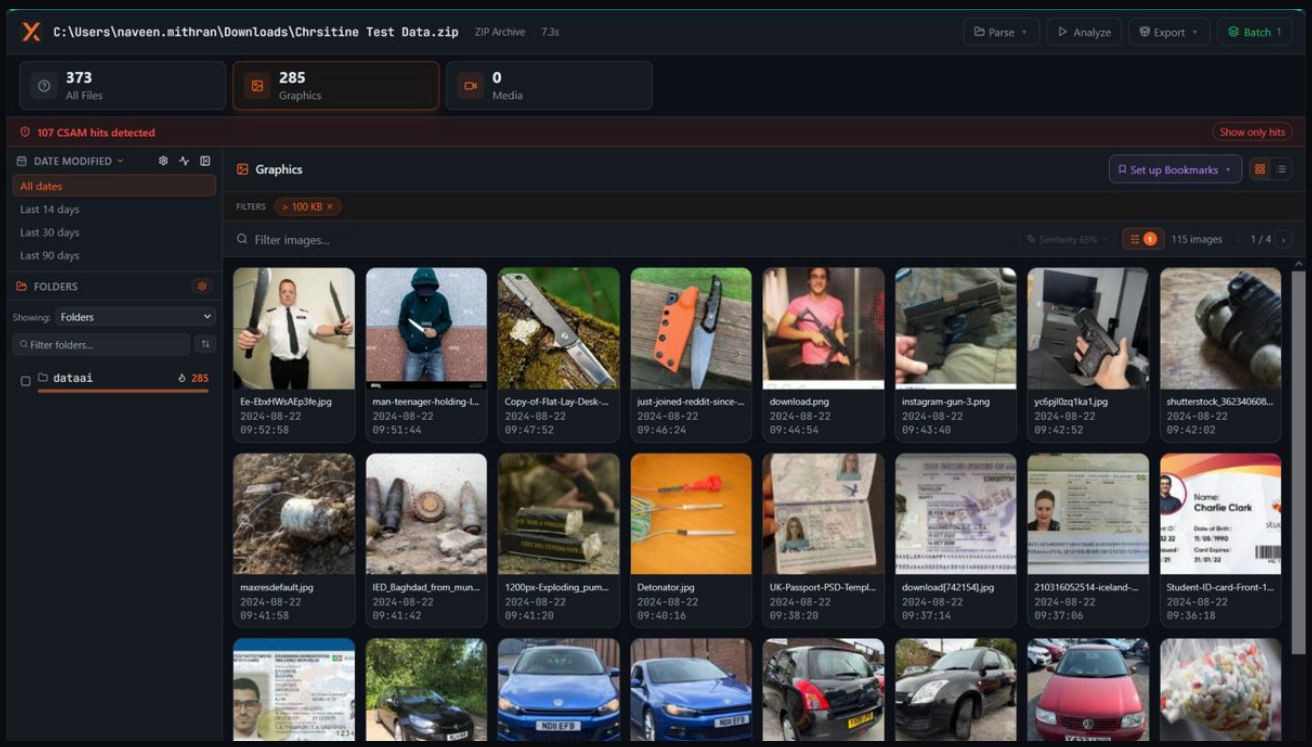

Toggle between thumbnail grid and list view using the icons in the top-right toolbar. In list view, images are shown with their filename, modified date, size, and path. When the CSAM Check has been run and hash matches are found, a red alert banner appears at the top showing the total number of hits. Click Show only hits to instantly filter the grid to only the flagged images.

The CSAM alert banner only appears when the CAID hash database has been loaded and the CSAM Check analysis has been run. See the CAID Lookup Builder section for setup instructions.

Media

The Media tab shows all video files discovered in the evidence. Toggle between list view and thumbnail grid view using the icons in the top-right toolbar.

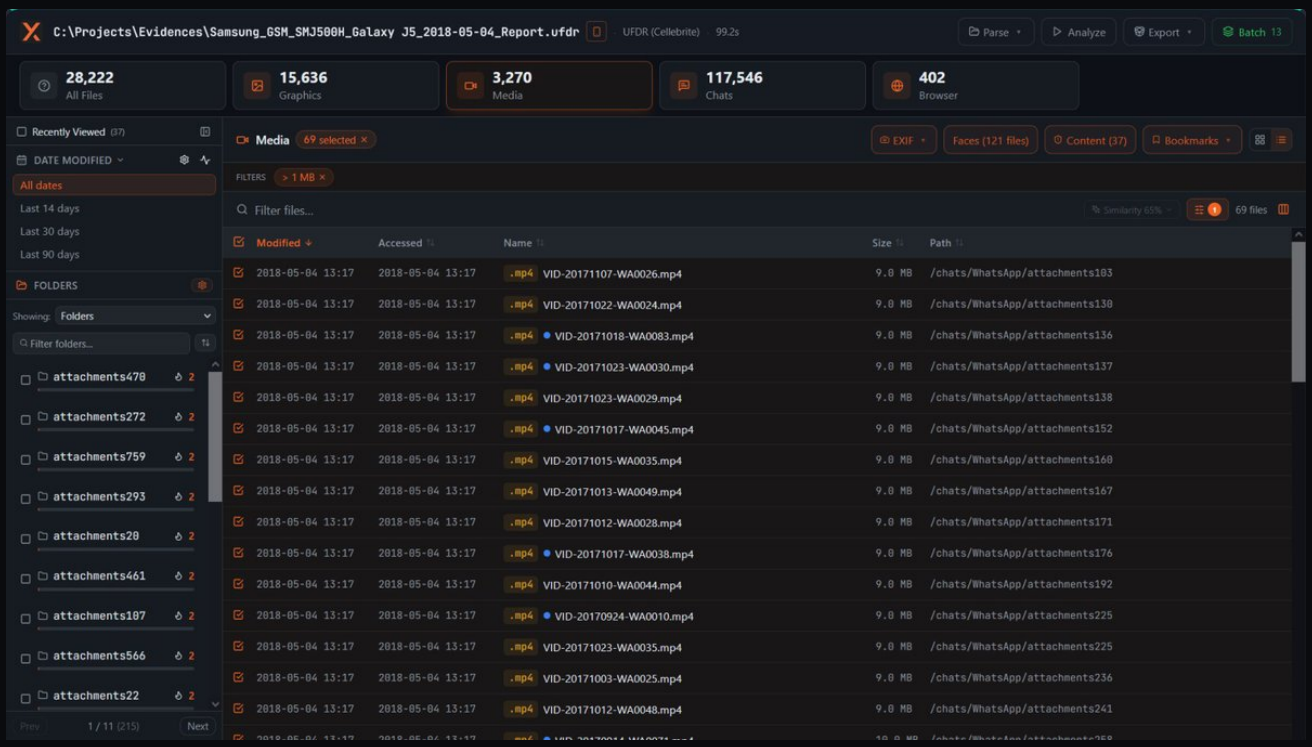

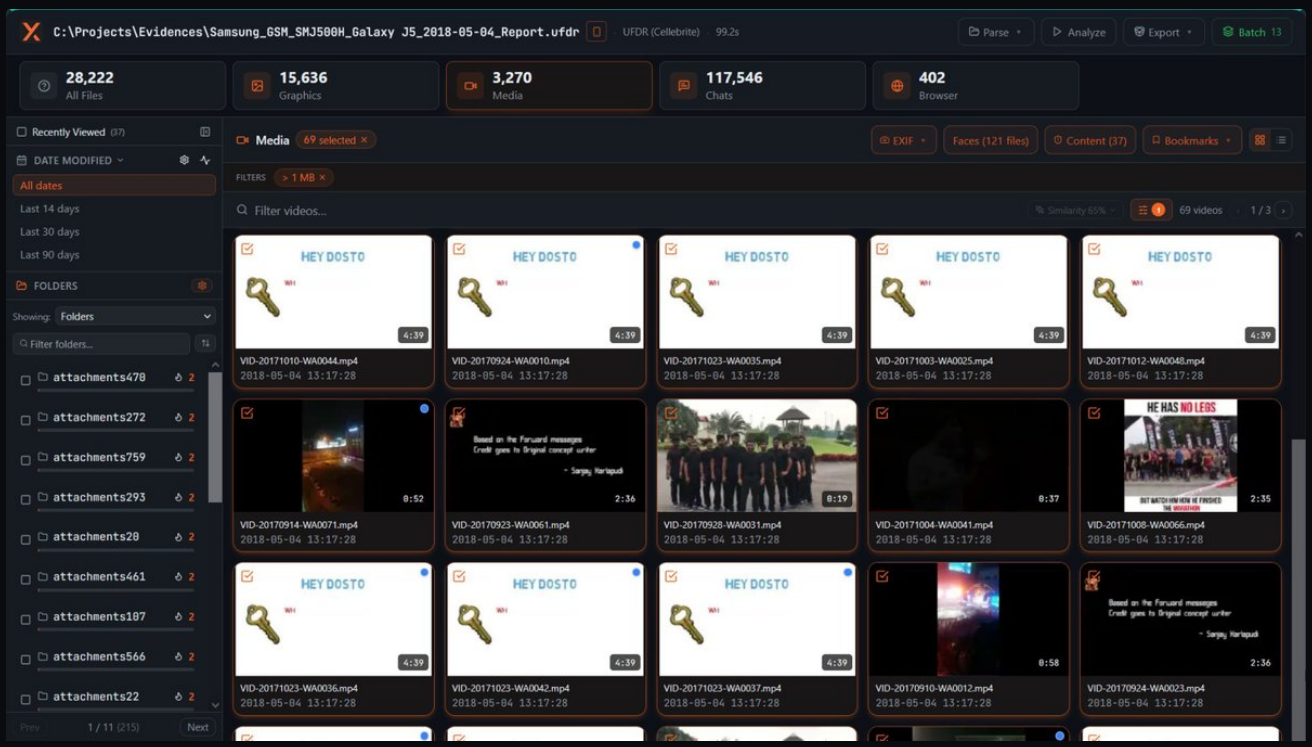

In grid view, each video card shows a preview frame and the video duration. In list view, videos are shown with filename, modified date, accessed date, size, and path.

Running AI analysis

Click the Analyze button in the top-right toolbar to open the analysis panel. From here you can select which AI operations to run on your selected files.

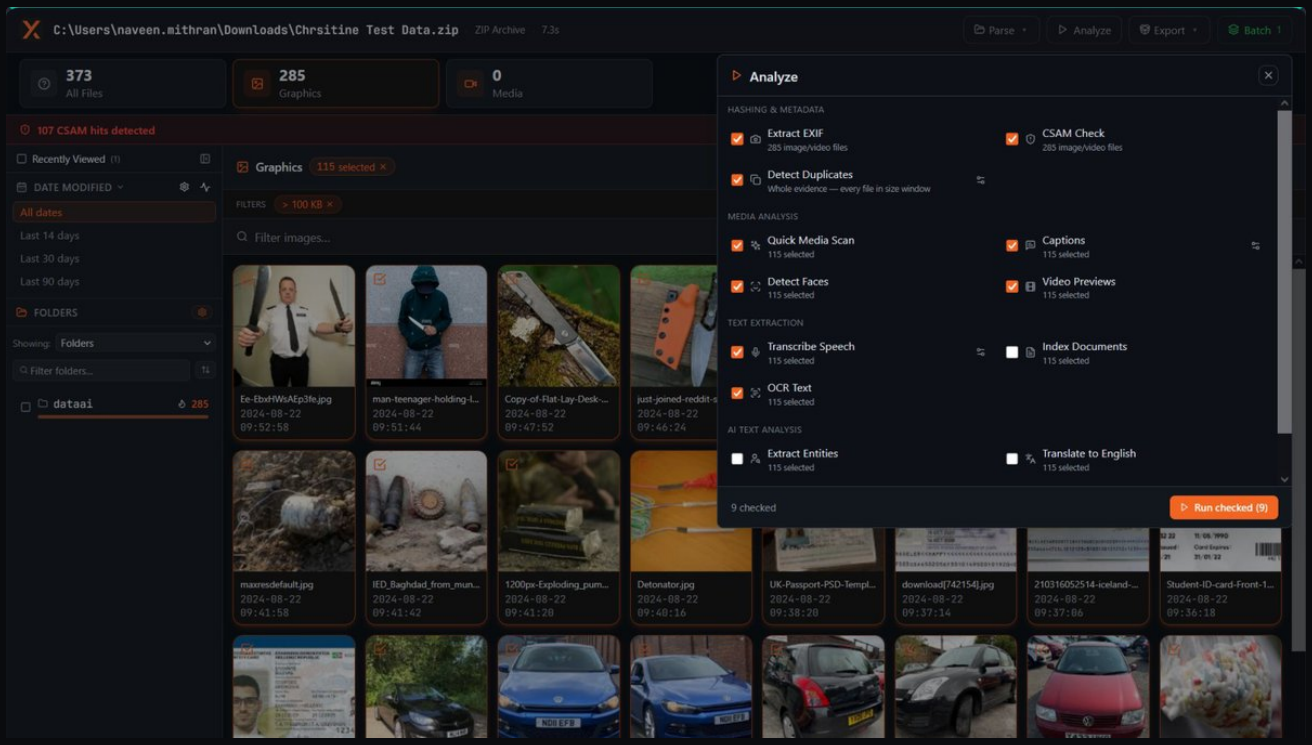

Each operation can be run independently or in combination. The analysis options are grouped into four categories: Hashing and Metadata, Media Analysis, Text Extraction, and AI Text Analysis. Select the files you want to analyze first, check the operations you want to run, and click Run checked to begin.

Each sub-section in this guide explains what each operation does and what it produces.

> You do not need to run all operations at once. Start with the ones most relevant to your investigation and run additional analysis later if needed.

CSAM Check

CSAM Check compares every image and video file against the CAID hash database loaded in the CAID Lookup Builder. Any file whose hash matches a known illicit image is flagged immediately.

Hashing And Metadata

Results appear as a red CSAM alert banner at the top of the Graphics tab, showing the total number of hits. Click Show only hits to filter the grid to only flagged files.

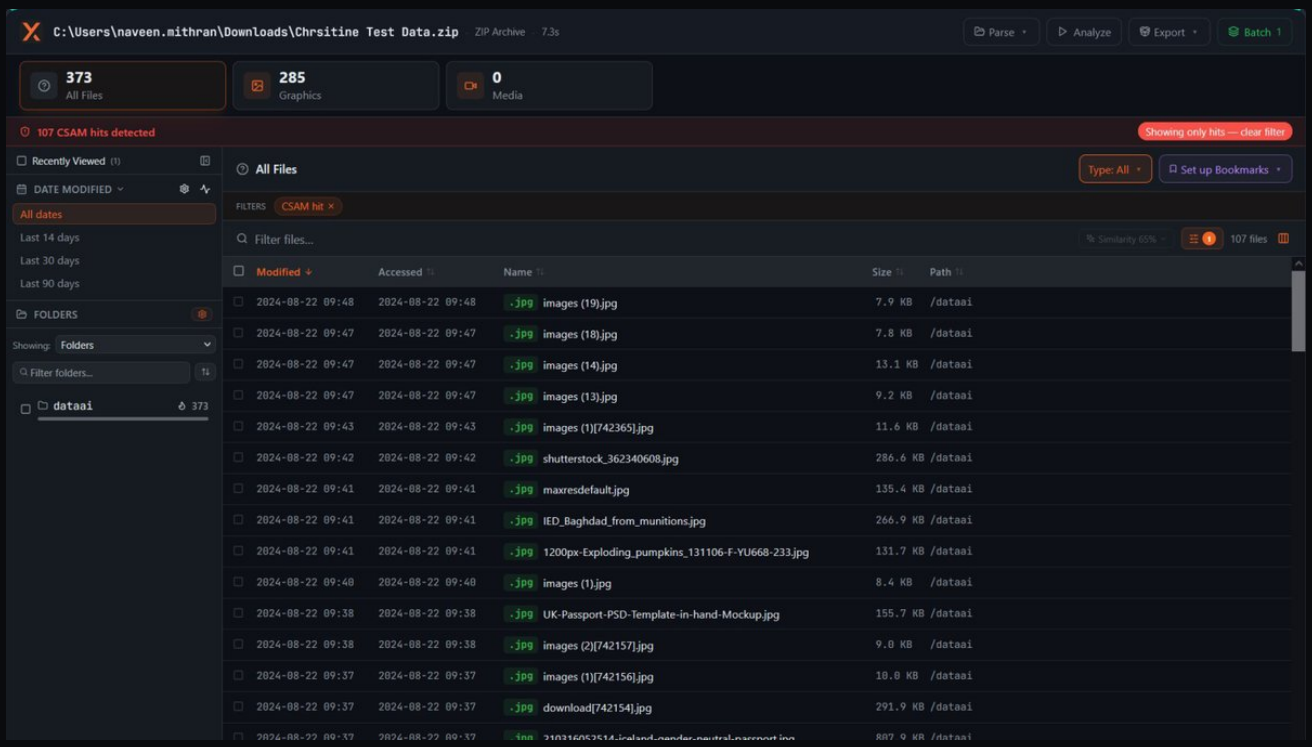

> This check requires the CAID hash database to be set up in advance. See the CAID Lookup Builder section for setup instructions.

Extract EXIF

Extract EXIF pulls camera metadata and GPS location data embedded in image and video files. Once extracted, two filter options become available in the EXIF toolbar button.

Hashing And Metadata

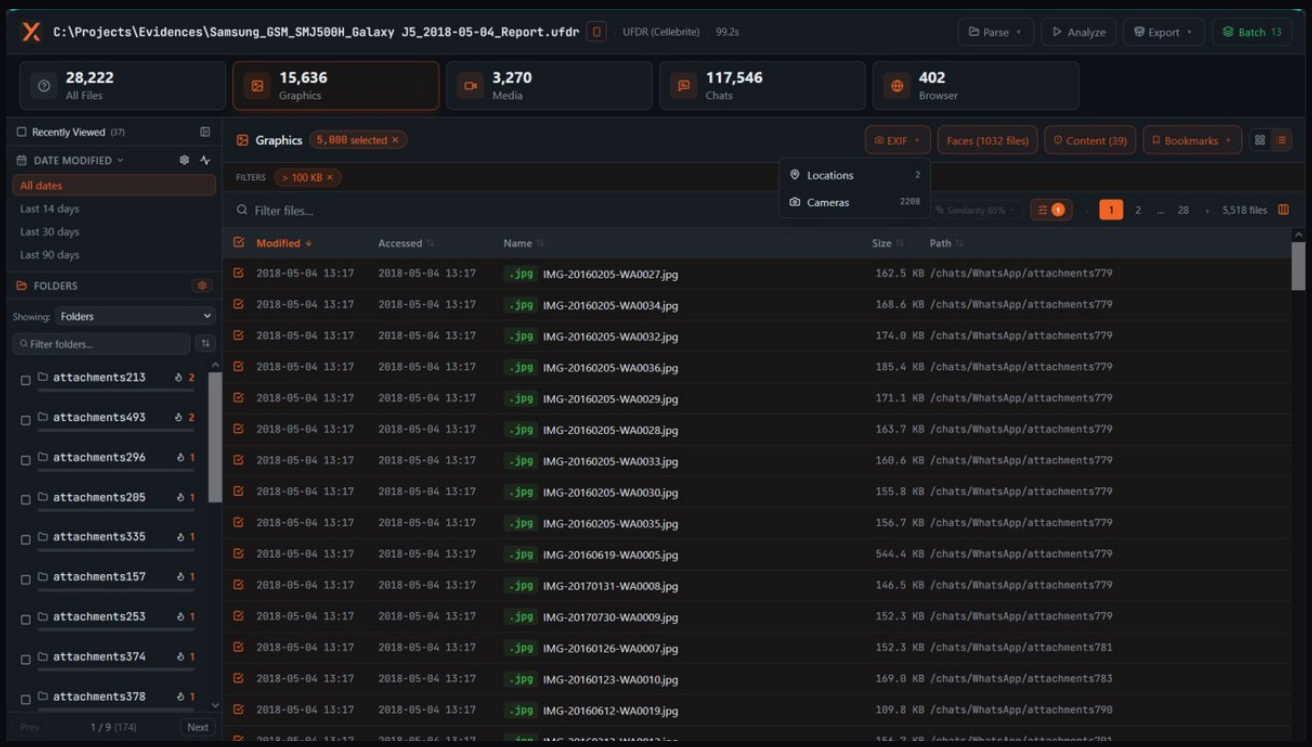

Clicking Locations opens a dialog listing every GPS location found across the evidence, with the city, state, country, and file count for each.

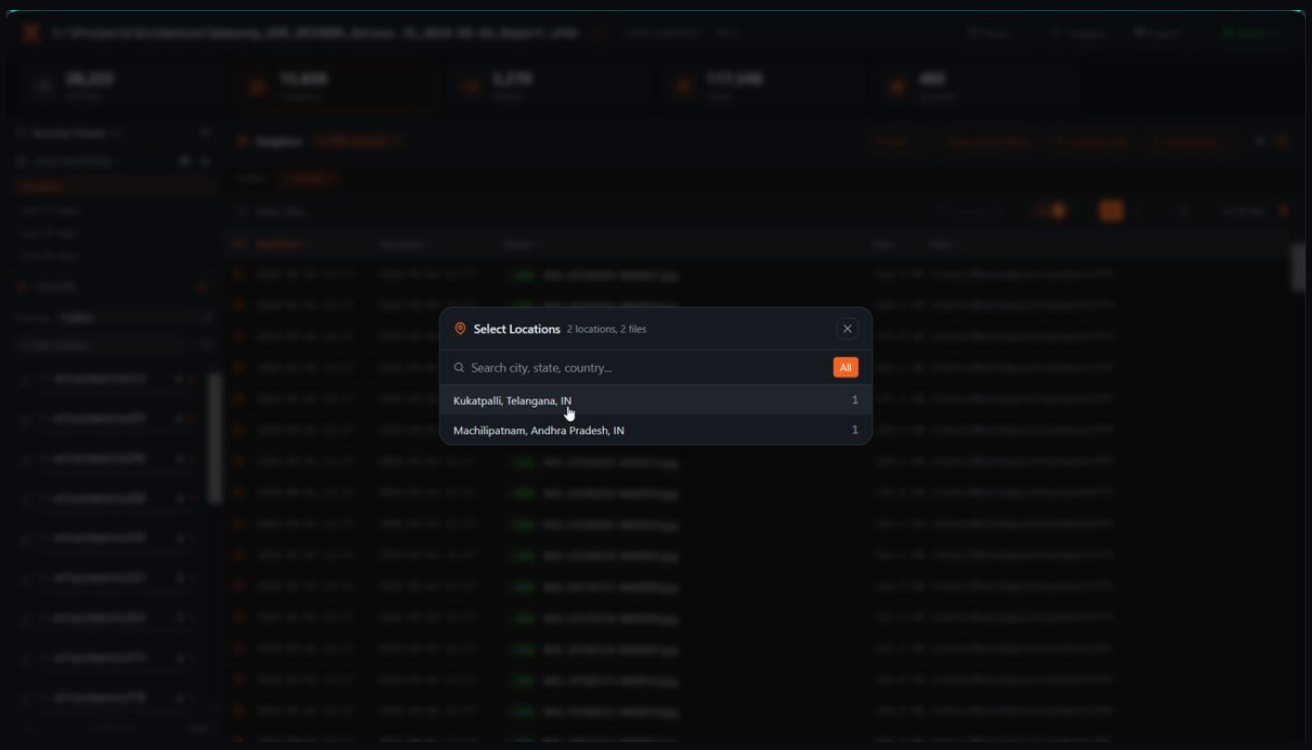

Select one or more locations to filter the image list to photos taken at that place. Cameras shows the count of images grouped by camera model.

Quick Media Scan

Quick Media Scan uses AI to scan and classify all images into investigatively relevant content categories.

Media Analysis

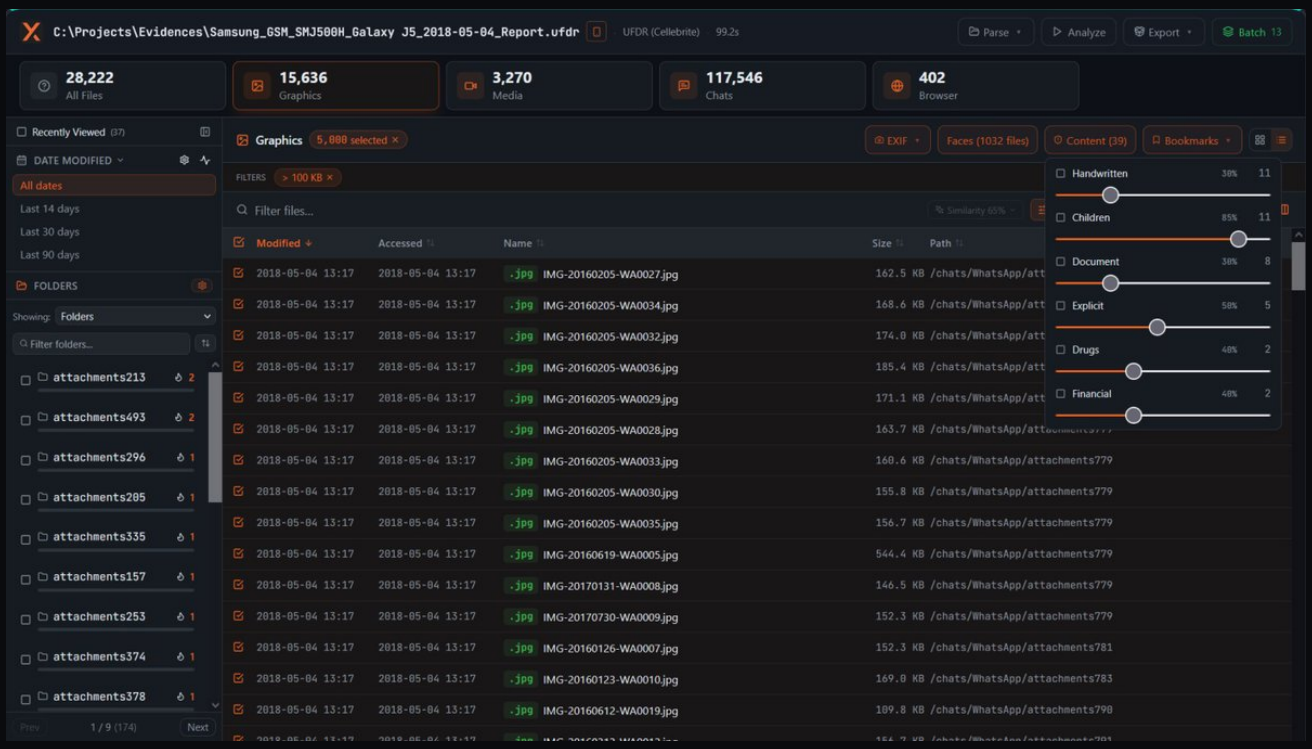

Handwritten Images containing handwritten text or notes. Children Images where the AI has detected the presence of children. Document Images of physical or printed documents. Explicit Images classified as containing explicit content. Drugs Images where the AI has detected drug-related content. Financial Images of financial documents, receipts, or currency.

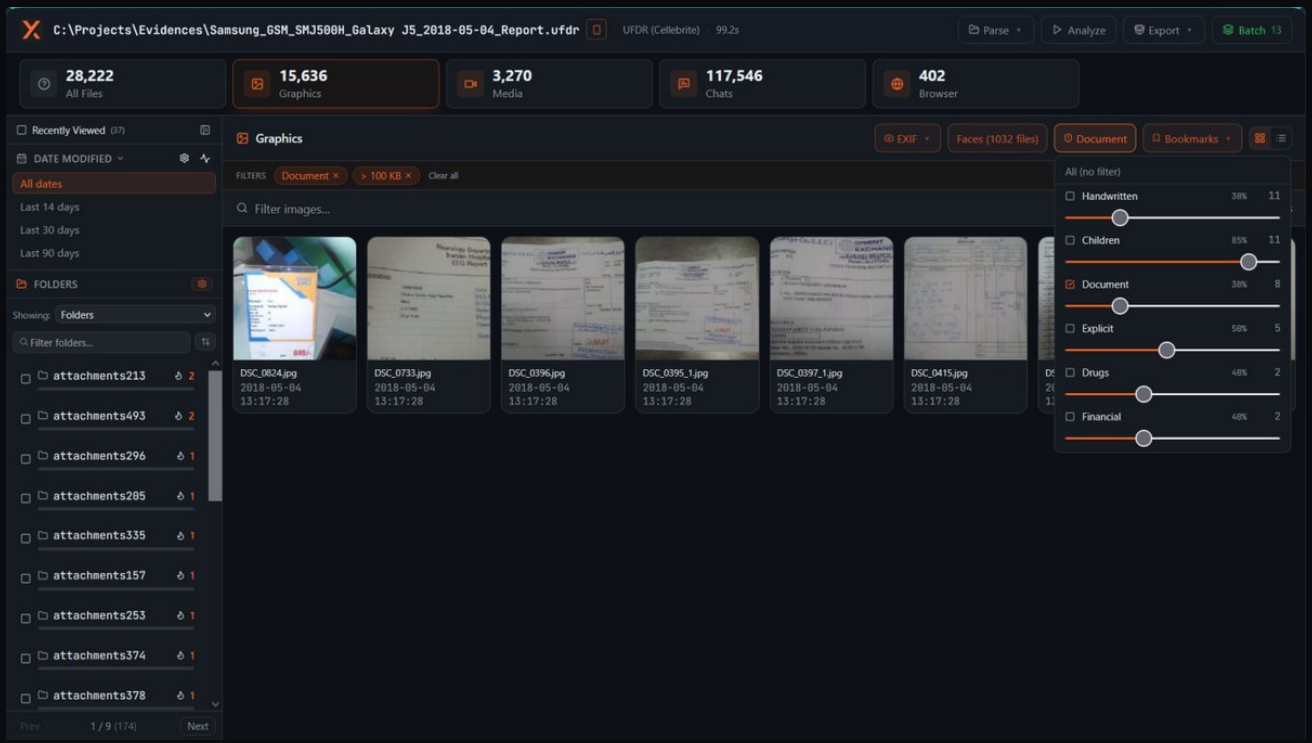

Each category has a confidence threshold slider. Drag the slider to set the minimum confidence level, then click Apply to filter the image grid.

Captions

Captions generates plain-language descriptions of what is in each image or video. For videos, the AI breaks the video into frames at regular time intervals and generates a timestamped caption for each frame.

Media Analysis

Captions make both images and videos fully searchable by keyword. Type any word or phrase into the search bar and the tool returns all files whose caption contains that term.

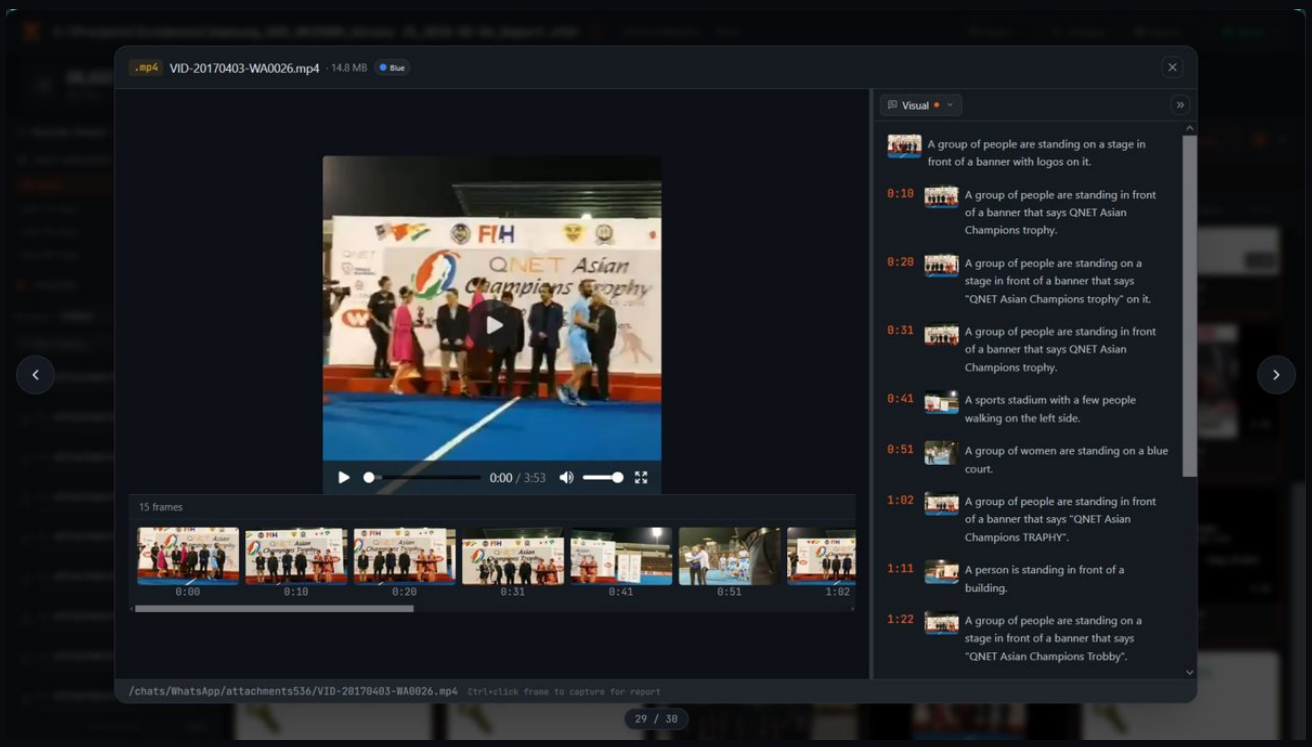

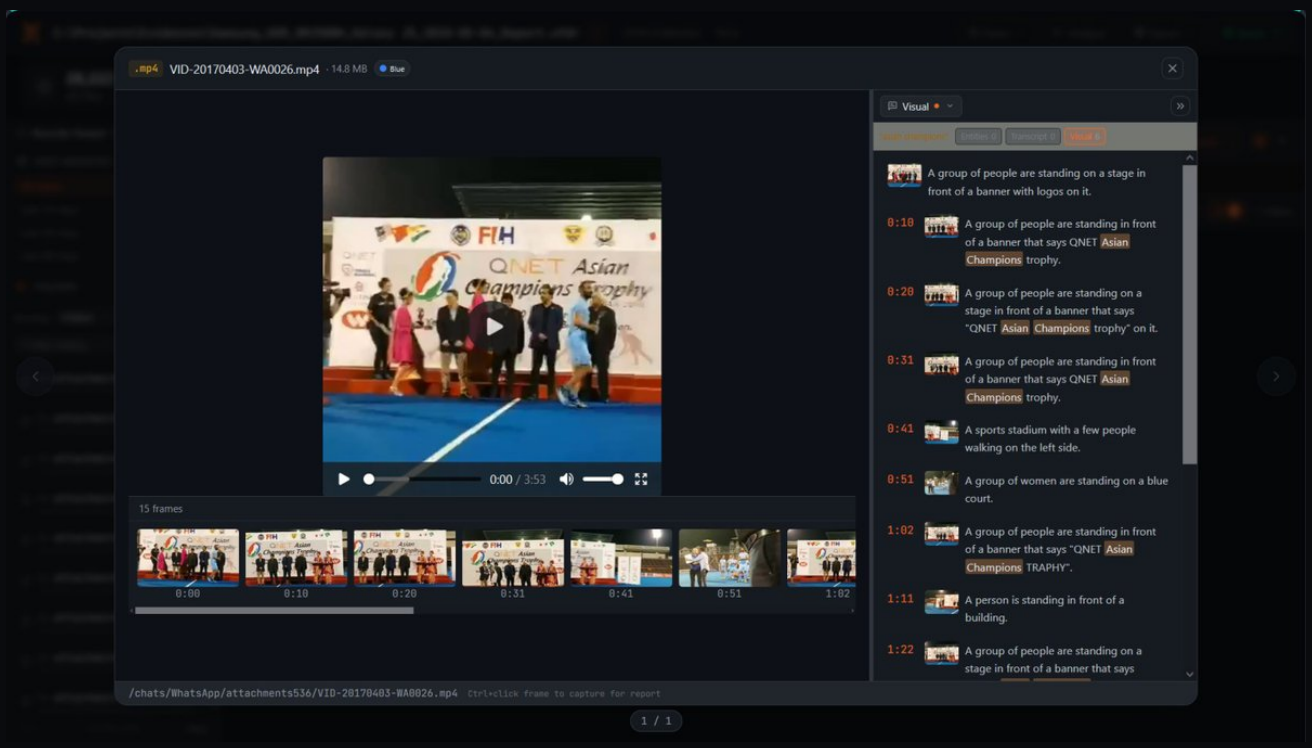

This is one of the most powerful features for video evidence. Hours of footage can be searched by a single keyword in seconds, without watching a single frame manually.

OCR Text

OCR (Optical Character Recognition) reads and extracts all text visible within an image, including text on physical documents, product labels, ID cards, signs, and whiteboards.

Text Extraction

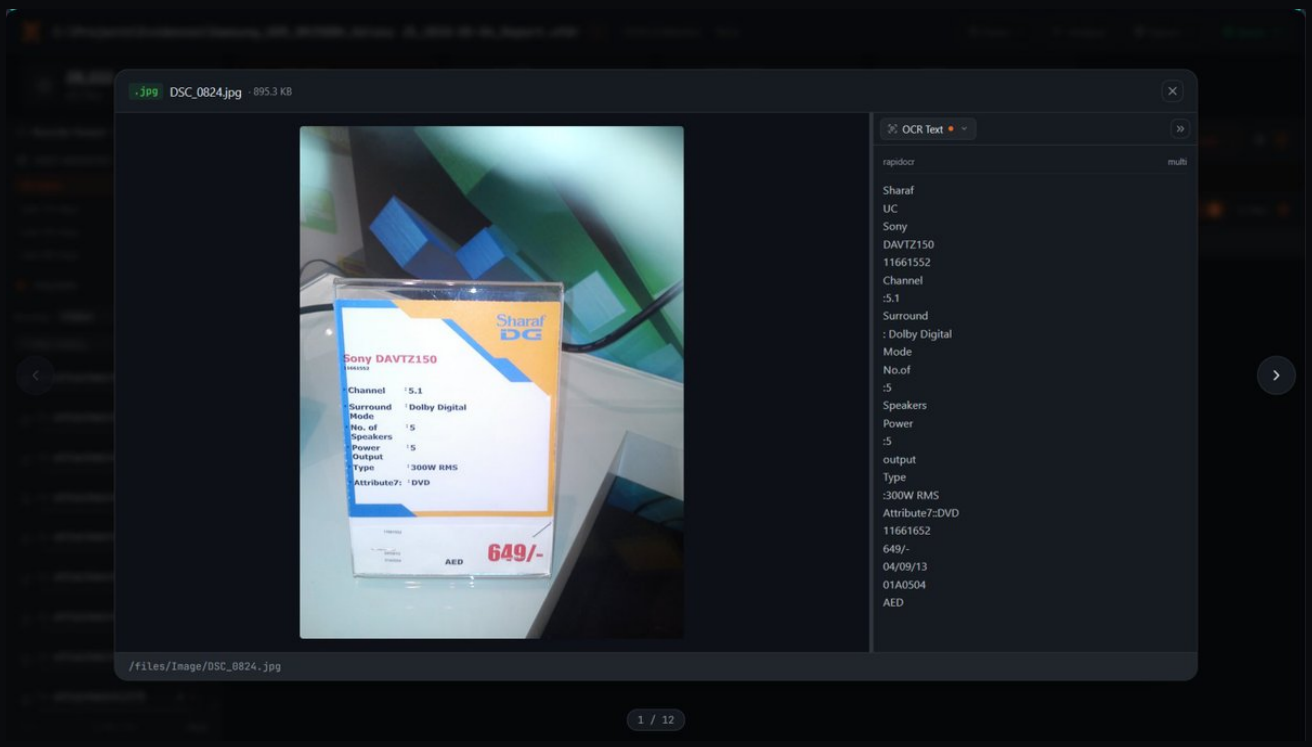

The extracted text is displayed in the OCR Text panel when you open a file, and is fully searchable across the entire job. This means investigators can search for a name, number, or phrase and find it even if it only appears as text within a photograph.

Transcribe Speech

Transcribe Speech converts spoken audio from video and audio files into timestamped text transcripts.

Text Extraction

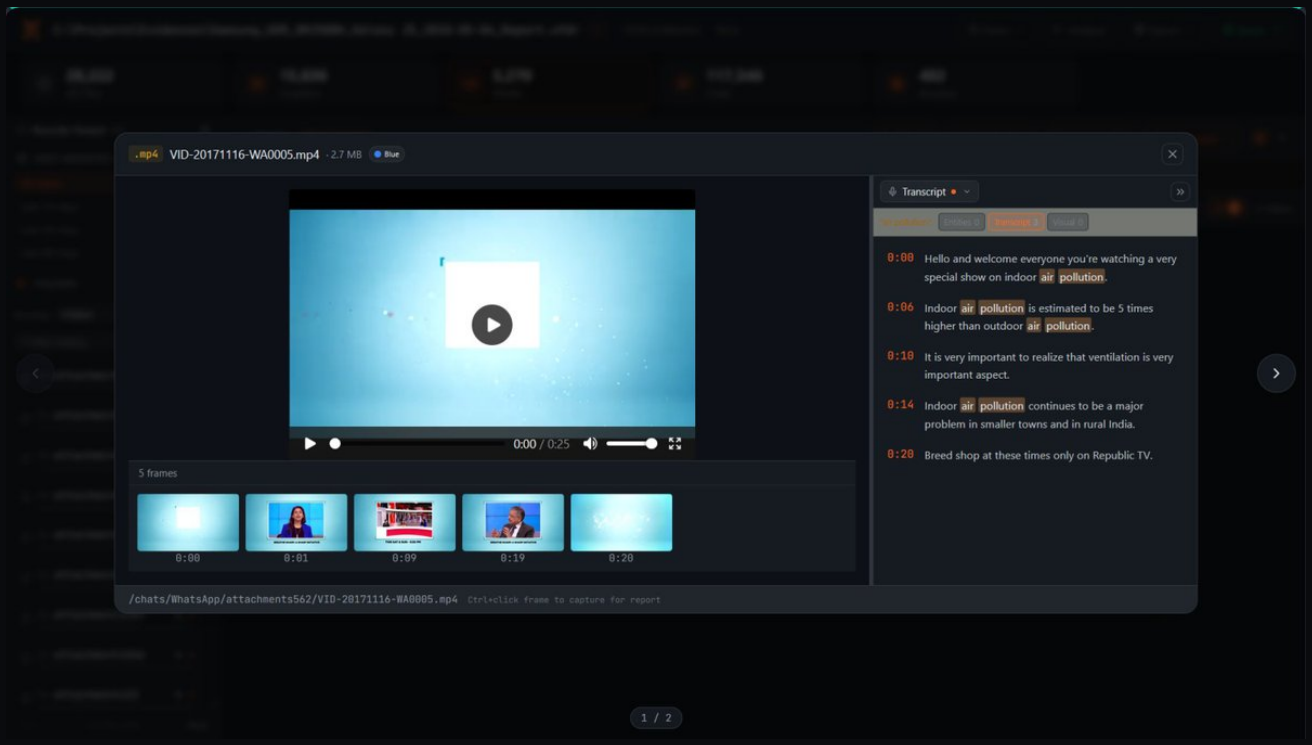

Each segment is tied to a specific timestamp so investigators can click any line in the transcript to jump directly to that moment in the recording. Transcripts are fully searchable, with matching keywords highlighted in the results.

> Hours of audio and video content can be searched by a single keyword in seconds. No manual listening required to locate a specific phrase or conversation.

Detect Faces

Detect Faces automatically detects every face across all images and videos and clusters them into individual person profiles.

Media Analysis

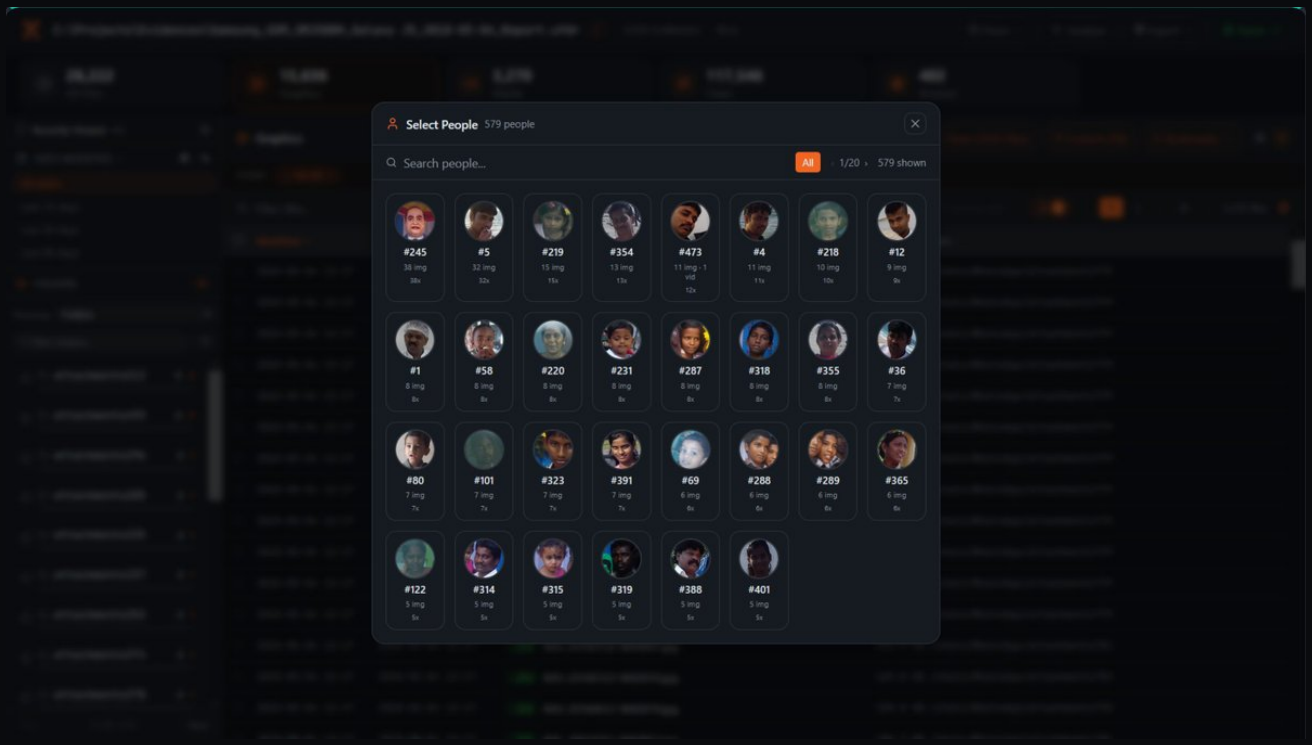

Each person is assigned a unique ID number and shown with a face thumbnail, the number of files they appear in, and how many times they appear in total. To filter the Graphics view to all images containing a specific person, click the Faces button in the toolbar, select one or more individuals, and click Apply.

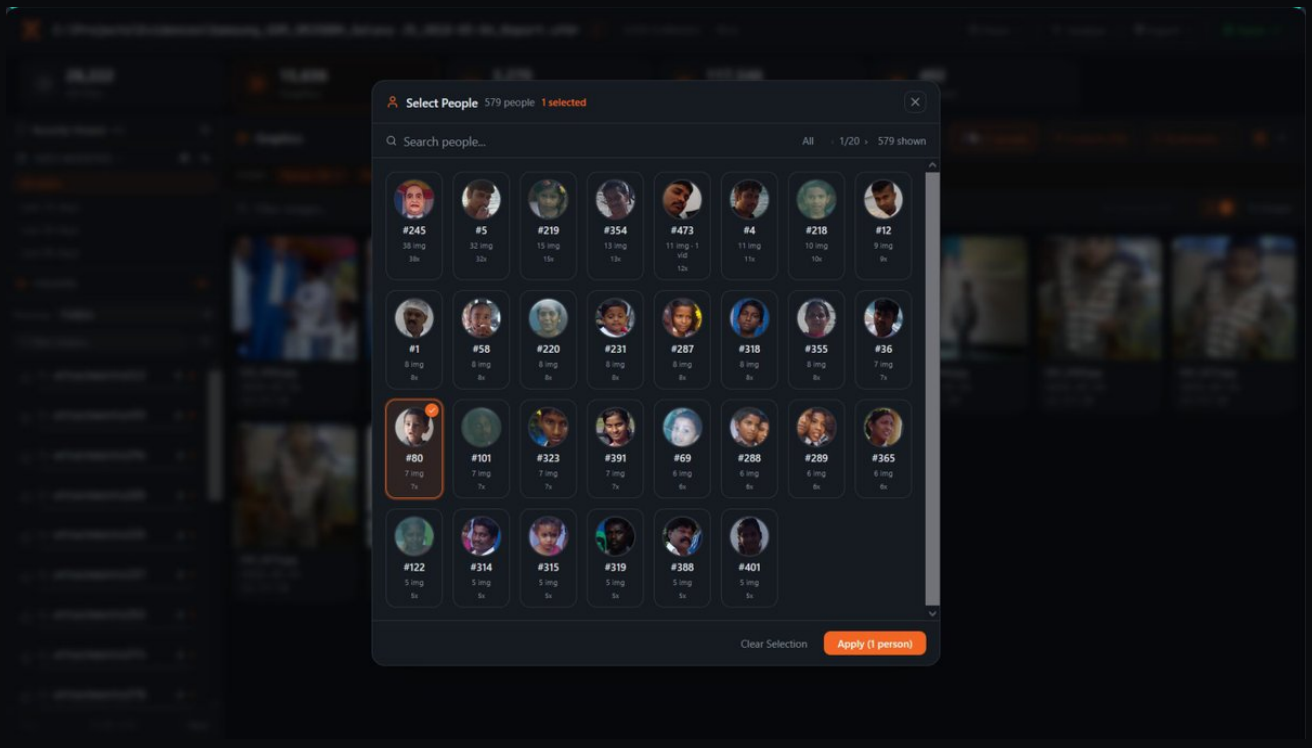

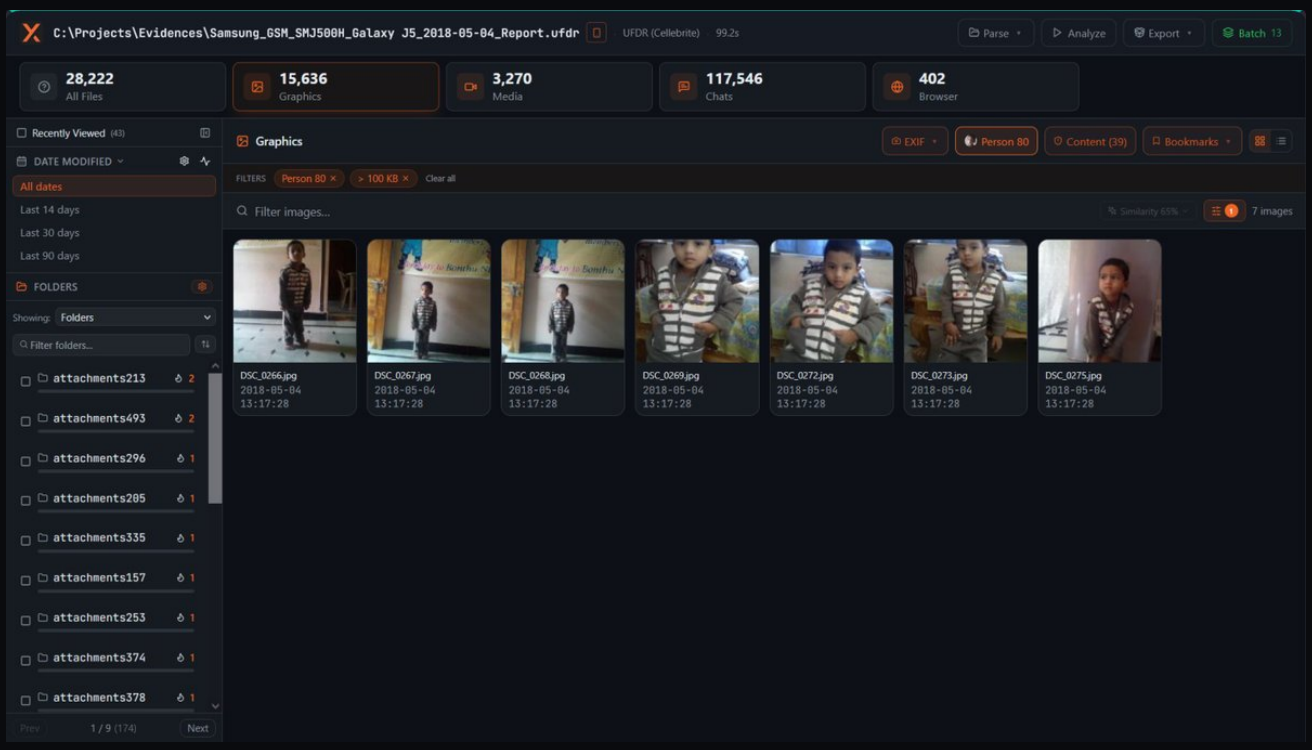

PART 6

DOCUMENT REVIEW

Indexing, searching, translating, and extracting entities from case documents

The Document Review section covers how to analyze and search through documents found in the evidence. Using AI-powered tools, investigators can index content for search, translate foreign language documents, extract named entities, and generate summaries. Why this matters: Evidence commonly contains hundreds of PDFs, Word documents, spreadsheets, and presentations.

ARMOURop lets investigators search across all document content using keywords or natural language, filter by entity names or locations, and read foreign-language documents in English instantly.

How To Access

From the All Files view, use the Type filter in the top-right toolbar to filter to office documents. Select the files you want to analyze and click Analyze.

Running document analysis

Select your document files, click Analyze, and choose which operations to run. The key analysis options for documents are in the Text Extraction and AI Text Analysis groups.

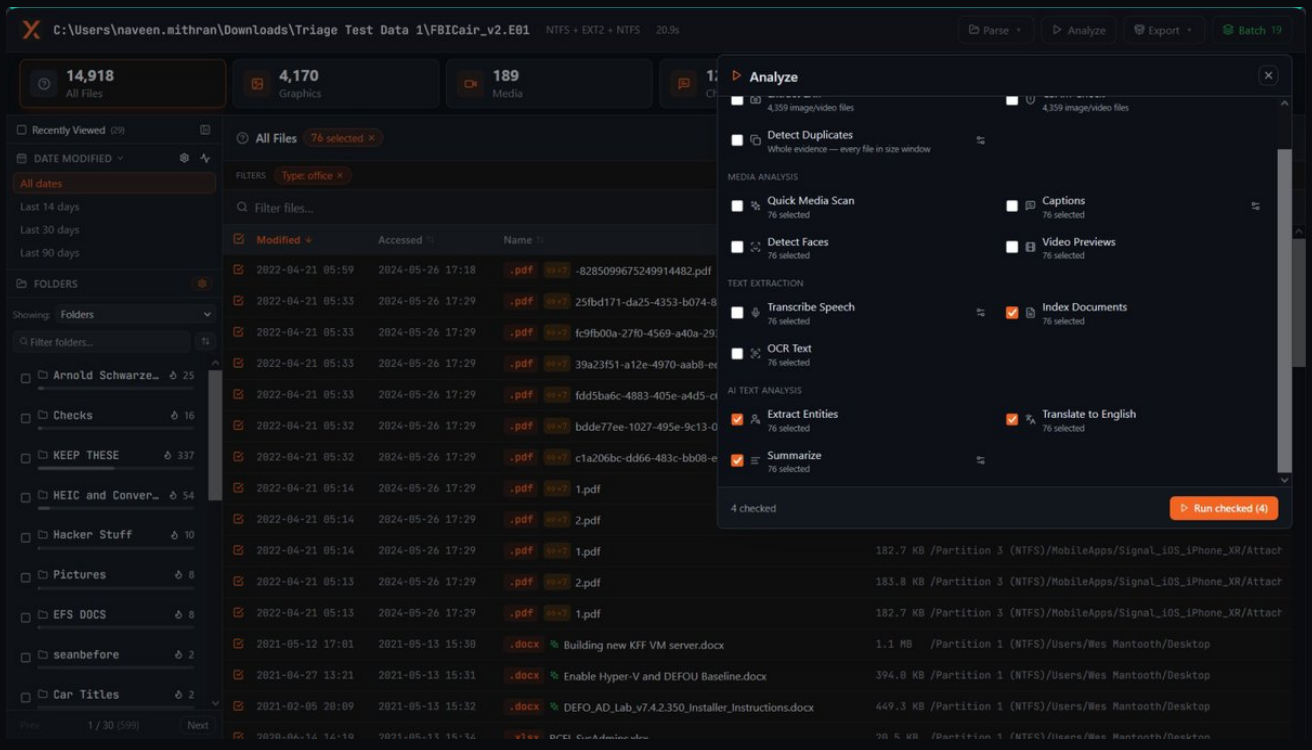

Check the operations you want to run and click Run checked. Each operation is explained in its own section of this guide.

Index Documents and search

Index Documents extracts and indexes the full text content of every selected document. Once indexed, documents become searchable by exact keyword and by AI-powered semantic meaning.

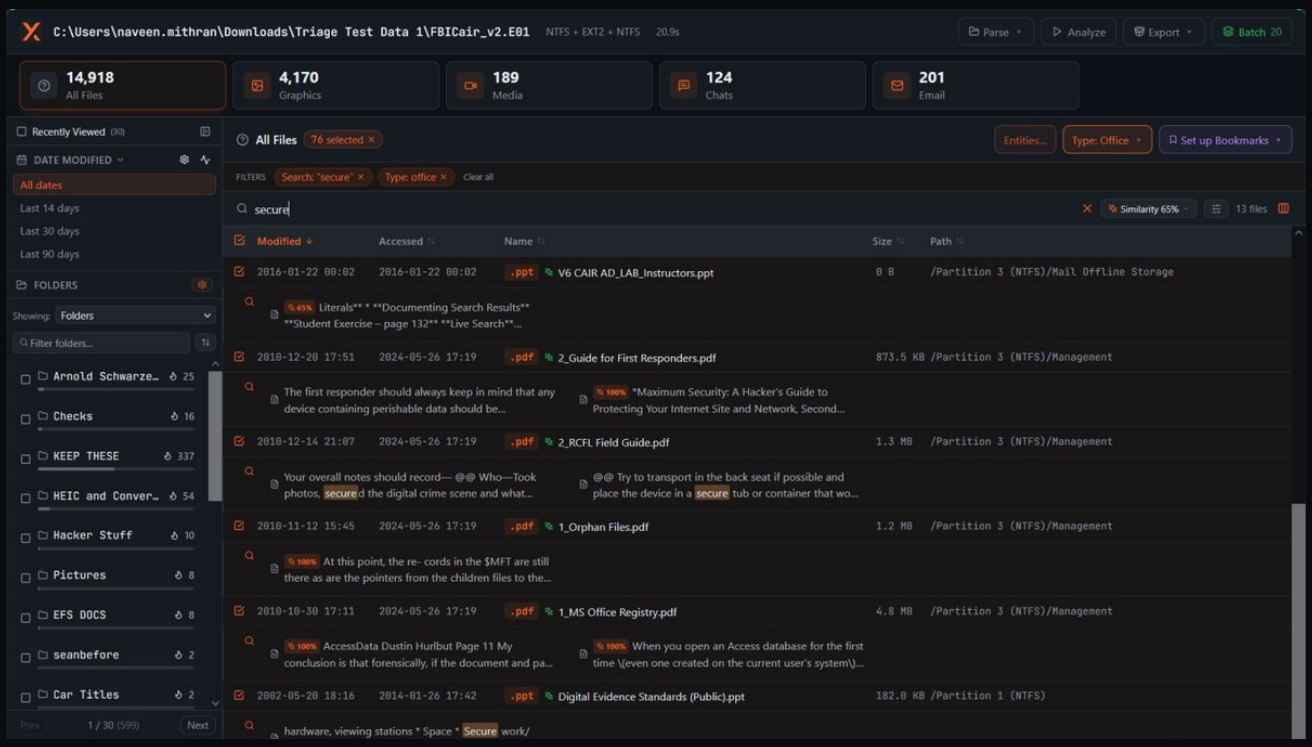

Each result shows a snippet of the matching content. Results fall into two types: Exact match The search term appears verbatim in the document. The matching word is highlighted in orange within the snippet. Semantic match The document content is semantically related to the search term, even if the exact word does not appear.

A confidence percentage badge shows how closely the content matches the meaning of the search. Adjusting the similarity threshold

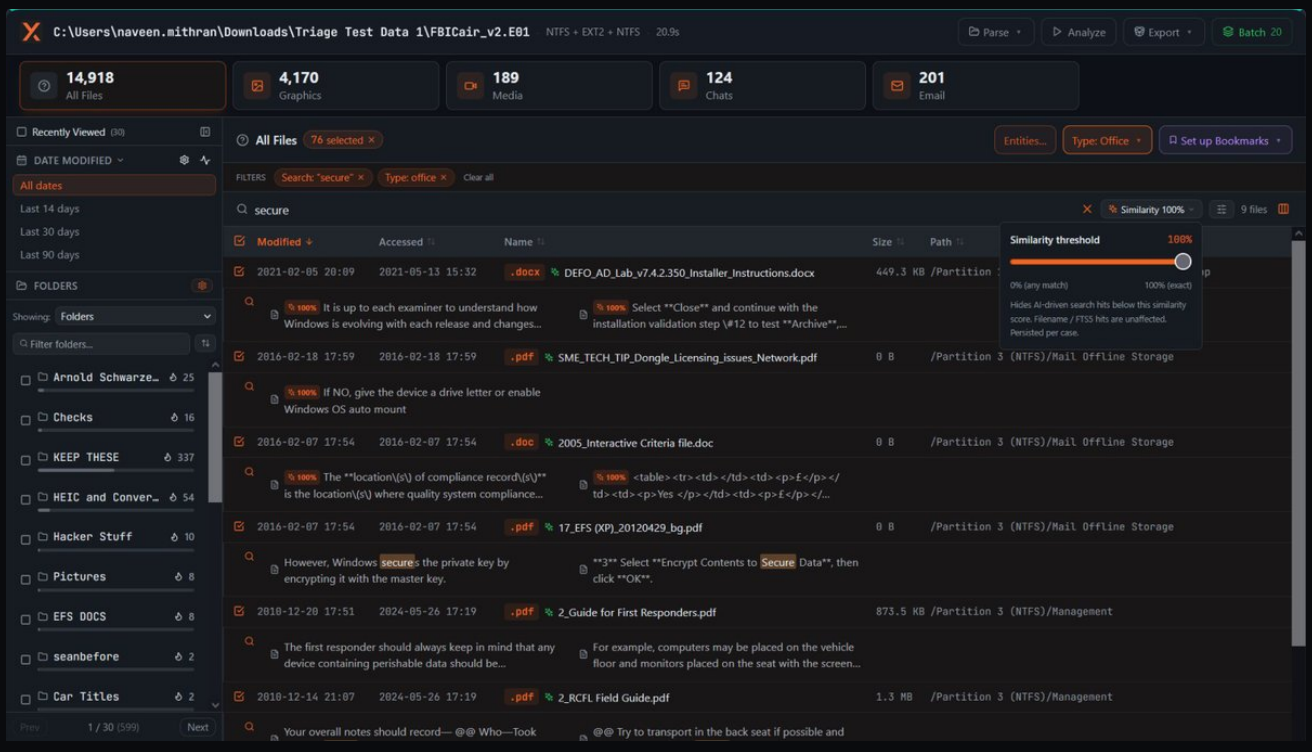

The Similarity threshold slider controls how strictly results must match. At 0% all semantically related results are included. At 100% only exact or near-exact matches are shown. The setting is saved per case.

> Use a lower threshold to cast a wide net at the start of an investigation, then raise it to narrow results once you know what you are looking for.

Translate to English

The AI automatically detects the language of each document and generates a full English translation. The original document is shown on the left, and the translation appears in the Translation panel on the right.

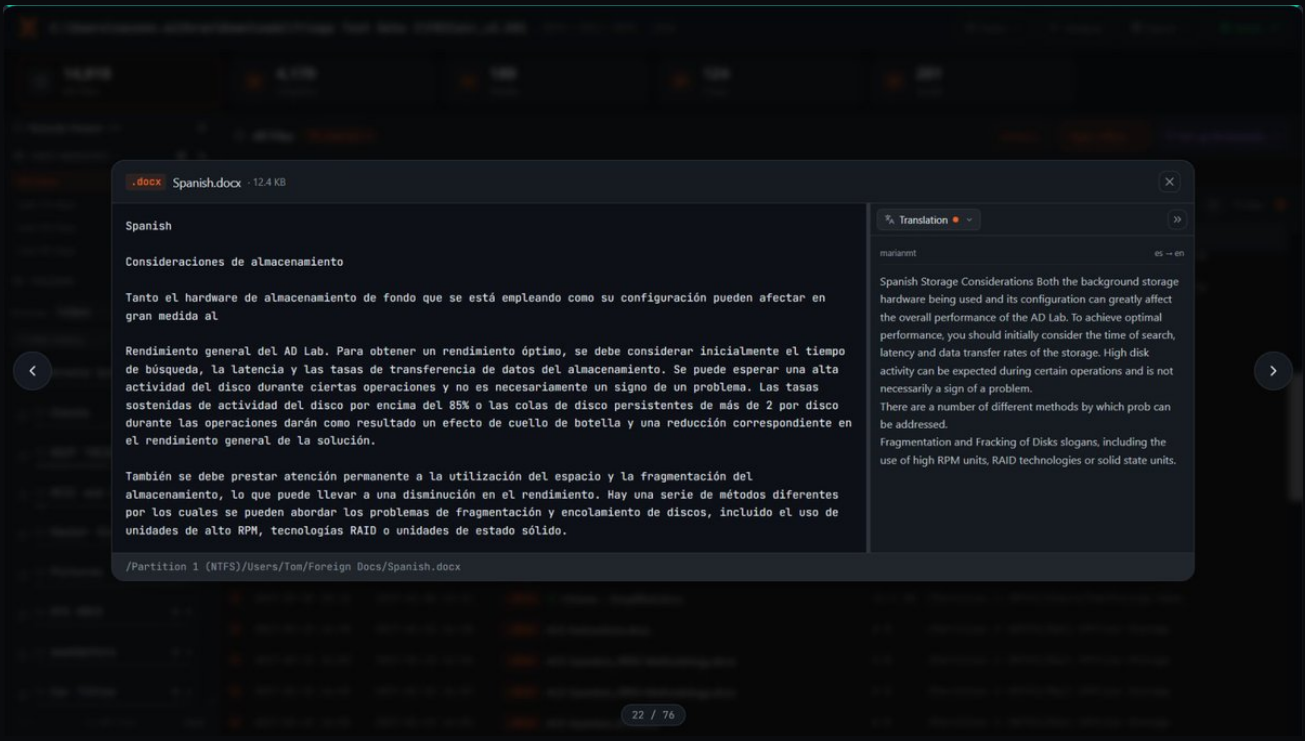

The language pair and translation model are shown at the top of the panel (e.g. es to en using marianmt).

> Translation runs entirely on-device. No document content is sent to any external service, maintaining chain of custody and compliance with evidence handling requirements.

Extract Entities

The AI reads every indexed document and extracts named entities such as organizations, locations, URLs, and vehicle plates. Once extracted, investigators can filter the entire document list by any entity value.

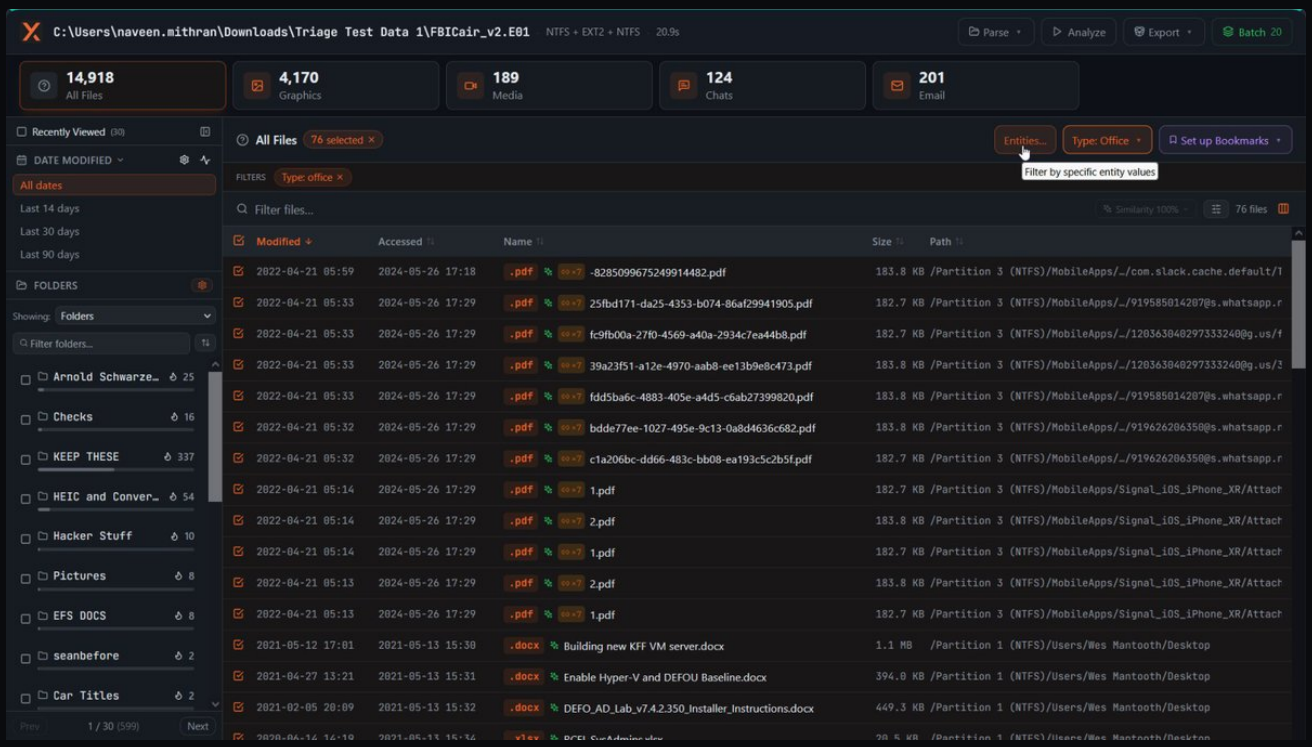

Click the Entities button in the top-right toolbar to open the Filter by Entity dialog.

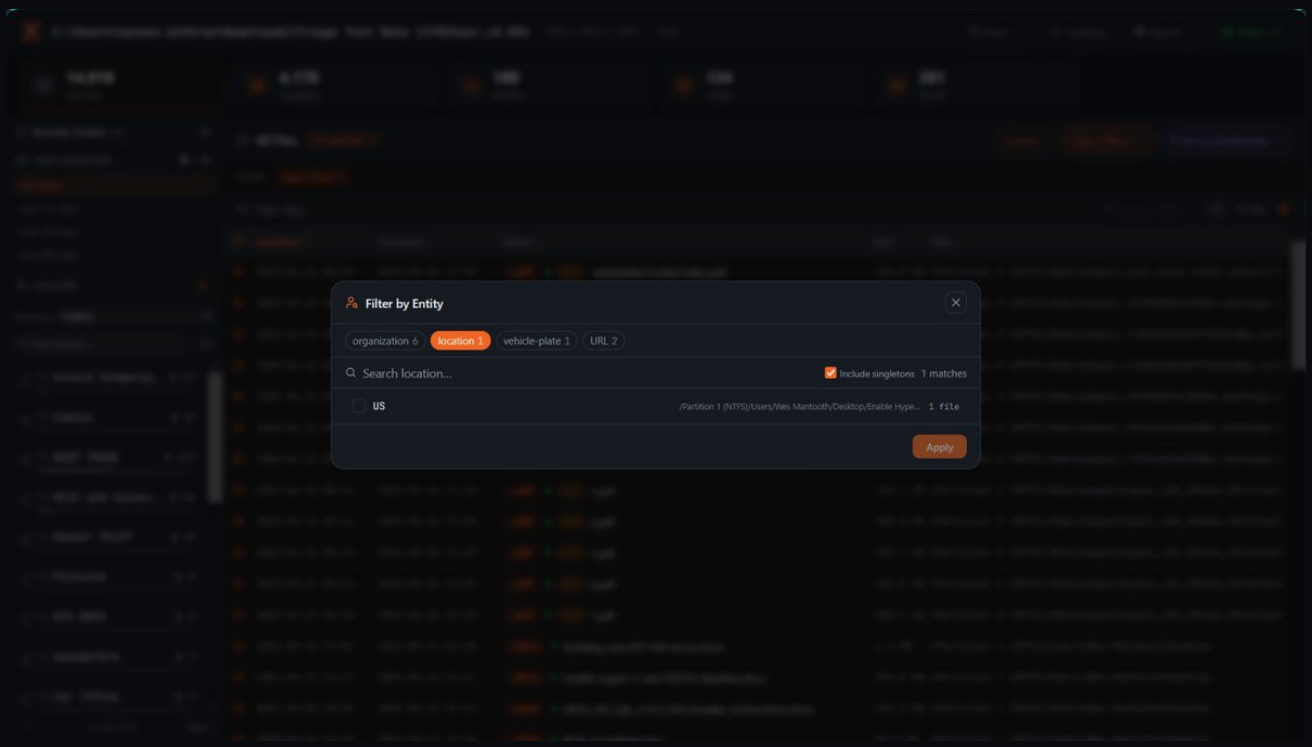

The dialog organizes extracted entities into tabs by type: organization, location, vehicle-plate, and URL. Select a value and click Apply to filter the document list.

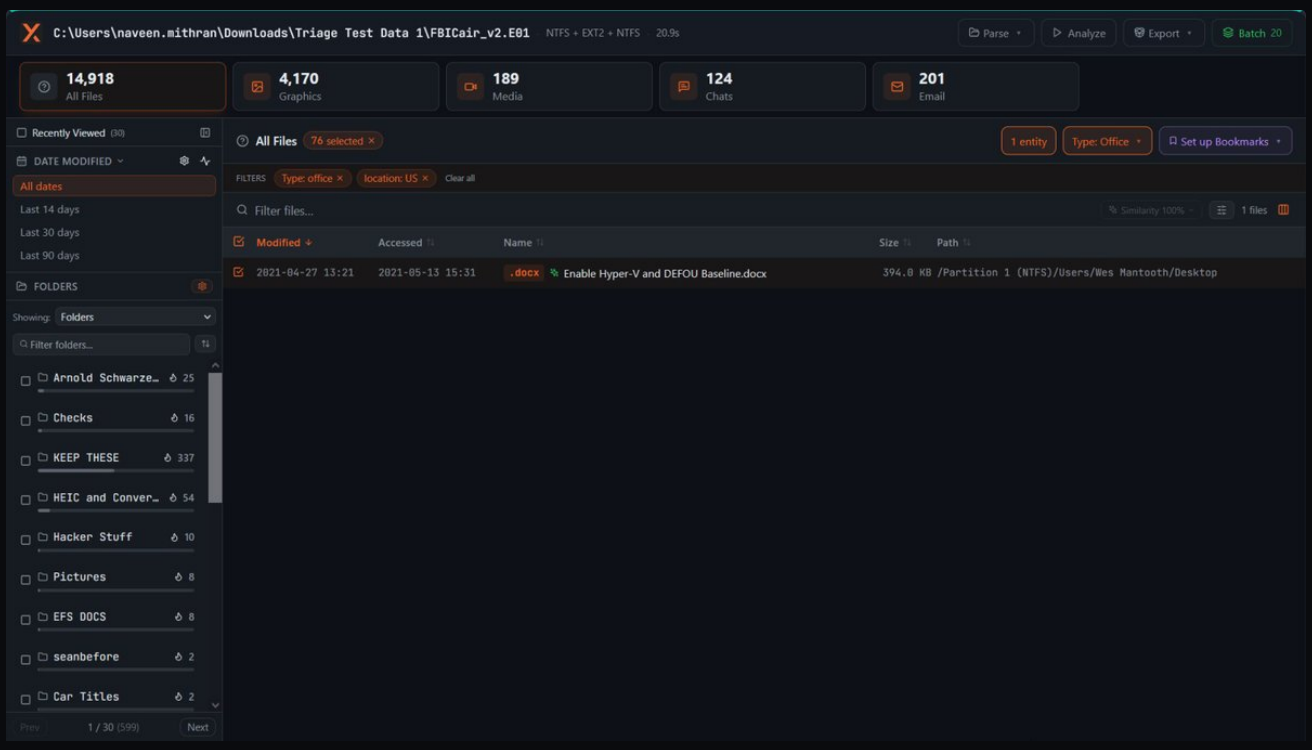

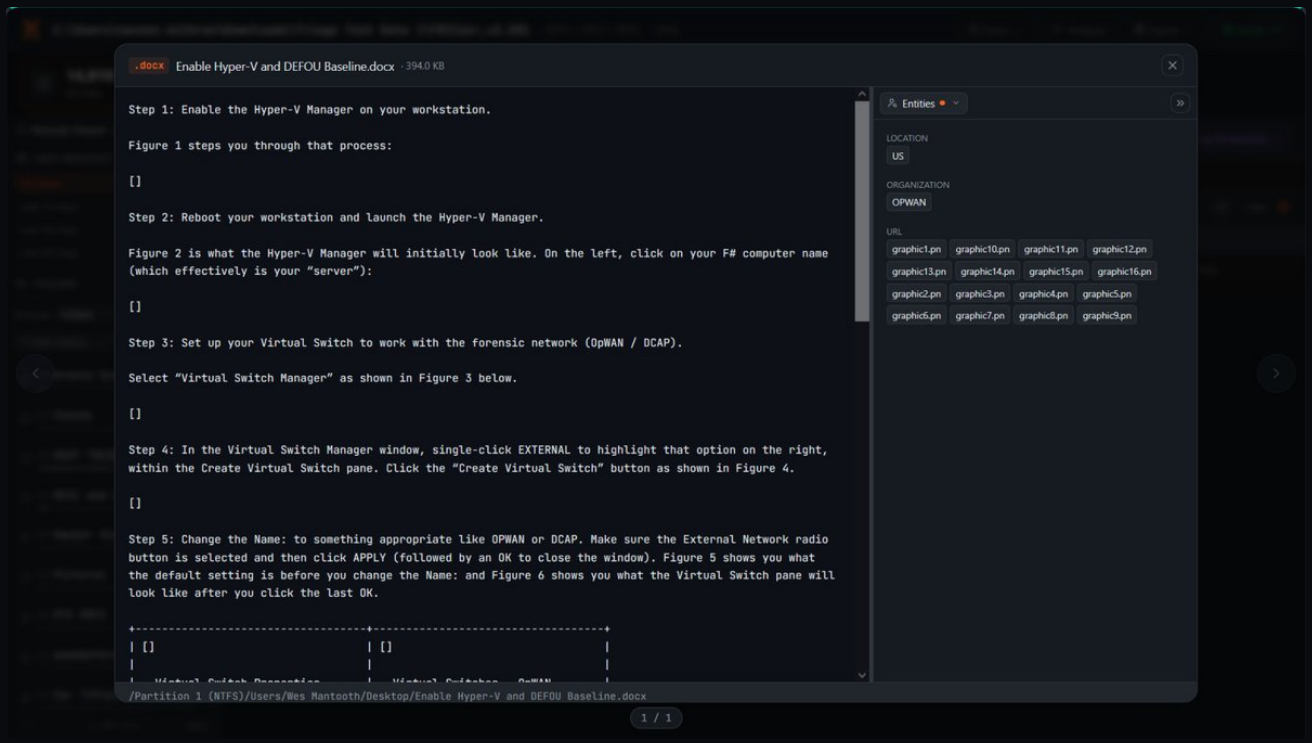

Use entity filtering to build a picture of connections across the evidence. Filtering by an organization name surfaces every document that mentions that organization, regardless of which folder it came from.

PART 7

EMAIL REVIEW

Parsing, reading, searching, and filtering PST and OST mailboxes

The Email Review section lets investigators read, search, and filter all email data found within the evidence. It supports PST and OST files from Outlook and Gmail exports. Why this matters: Email evidence can contain thousands of messages across multiple mailboxes.

ARMOURop consolidates all email sources into one searchable interface with filters for sender, domain, attachments, and folders.

How To Access

Click the Email tab in the file category summary bar at the top of the Triage Console. Email data must be parsed first before it can be reviewed.

Parsing email files

Before you can review email data, you need to parse the email files found in the evidence. Click the Parse button in the top-right toolbar and select Parse Email from the dropdown.

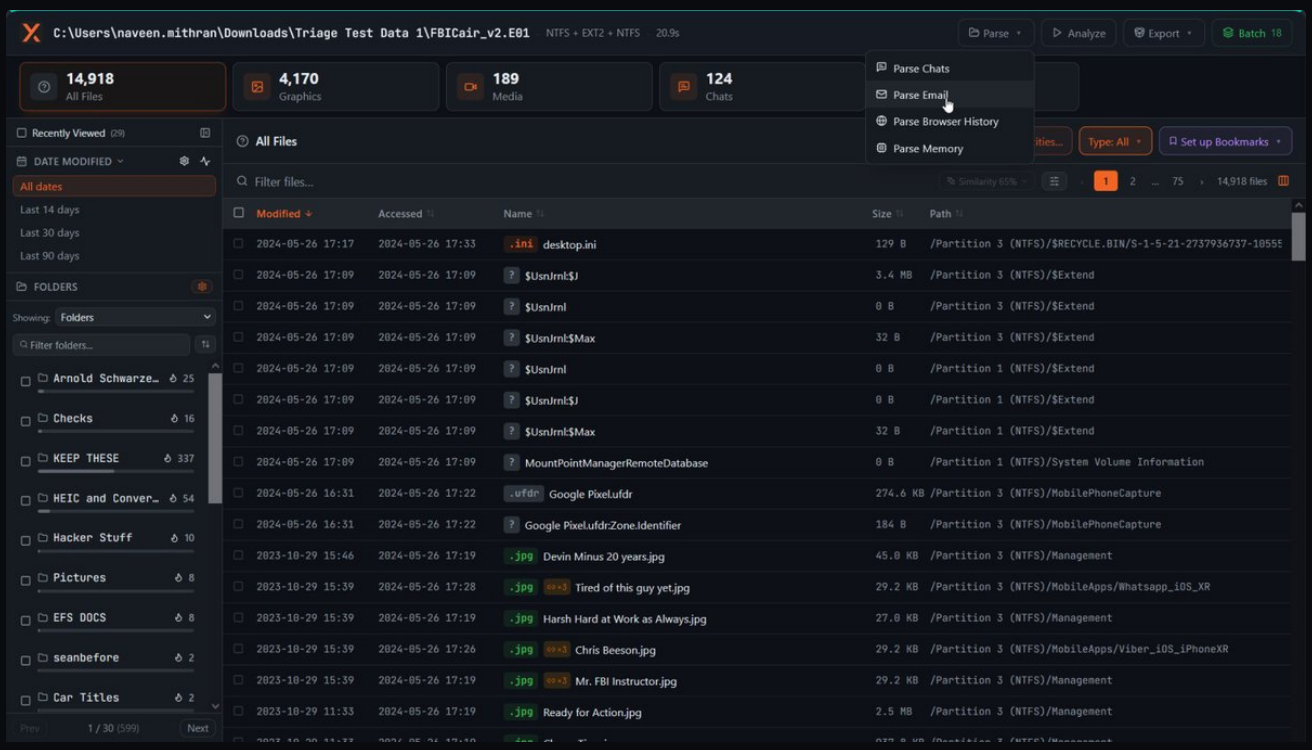

The Parse Email dialog lists all PST and OST files discovered in the evidence, showing the file name, path, format, and size. Files that have already been parsed are marked with a PARSED badge. Select the files you want to parse and click Parse Files.

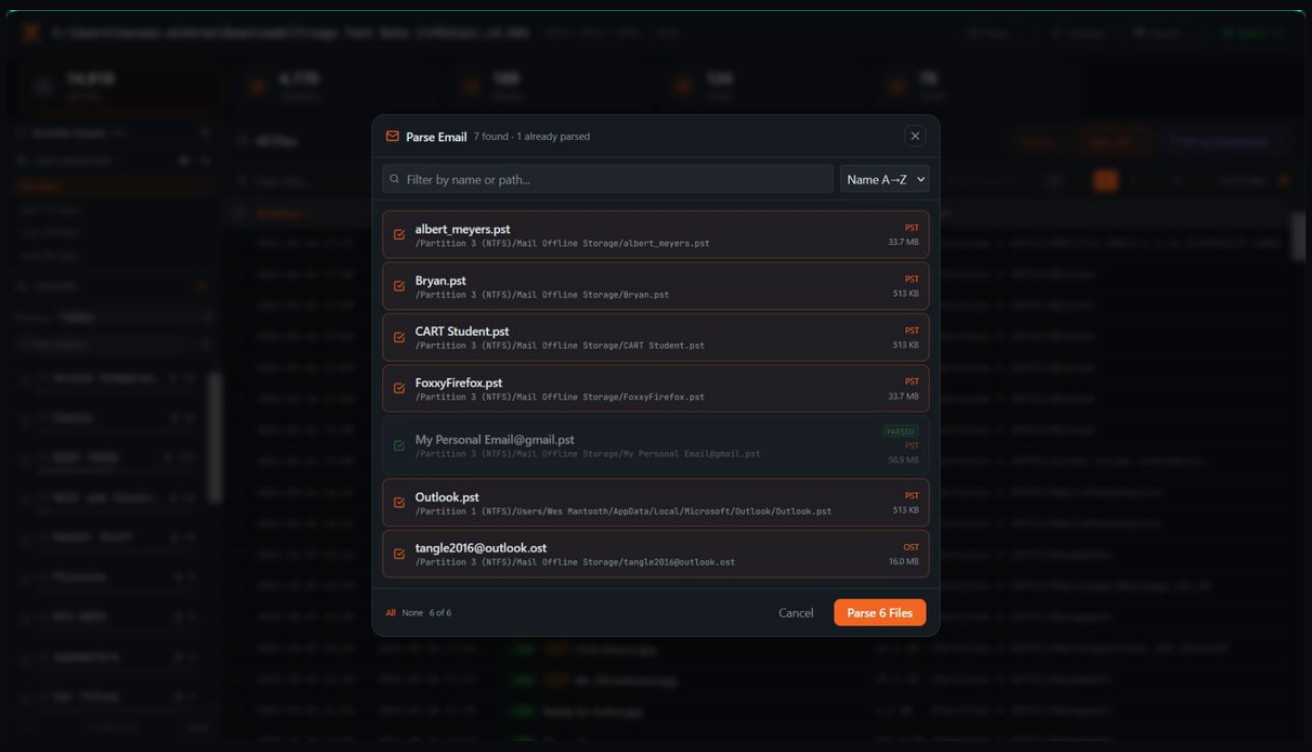

Parsing only needs to be done once per file. Files marked PARSED are skipped automatically if you run Parse Email again.

Reviewing emails

Once parsed, the Email tab offers two view modes: Conversation view, a three-panel layout for reading threaded mail in context, and Grid (table) view, a spreadsheet-style list for scanning and batch-selecting large volumes of email. Switch between them using the two icon buttons at the top right of the Email tab.

Conversation View

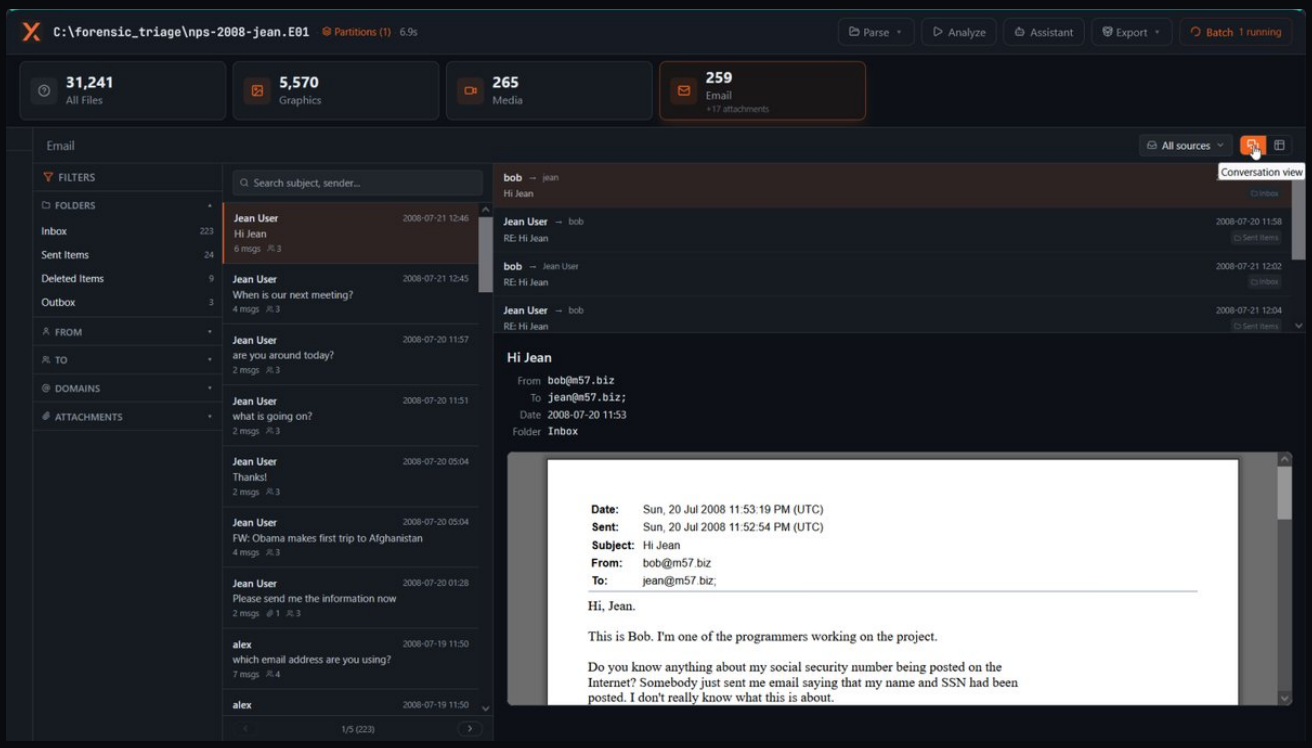

Folder navigation All email folders from the parsed files are shown with message counts: Inbox, Sent Items, Deleted Items, Outbox, and more. Click any folder to show only messages from that location. The FROM, TO, DOMAINS, and ATTACHMENTS filters sit below the folder list.

Thread list and email reader Each row in the center list shows the sender, subject, date, message count, and participant count. Click any thread to open it on the right, expanded into its individual messages with full headers, rendered body, and a page indicator at the bottom of the list.

Grid (Table) View

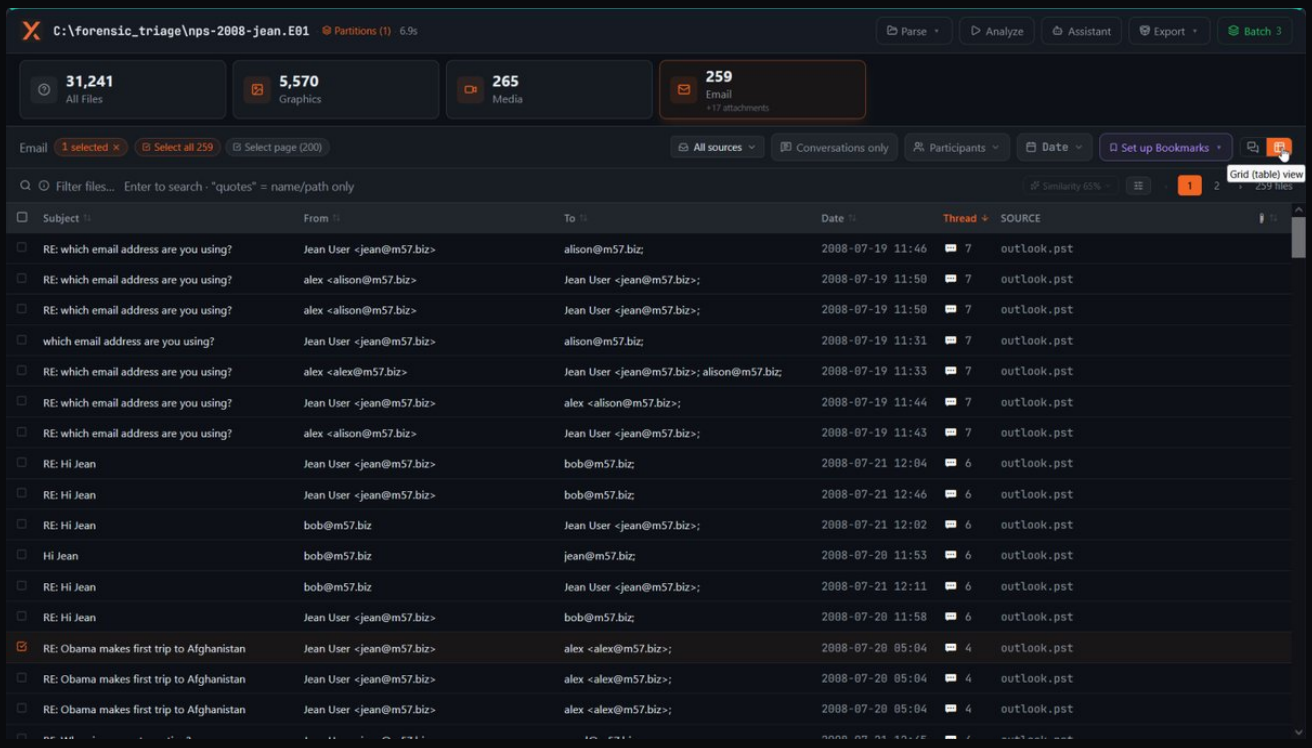

Grid view lists every email individually rather than grouped by thread, making it well suited to sorting, bulk selection, and export. Click any column header to sort by that column. Use the checkboxes to select individual emails, or Select all / Select page above the list to select in bulk, then use Set up Bookmarks to tag the selection.

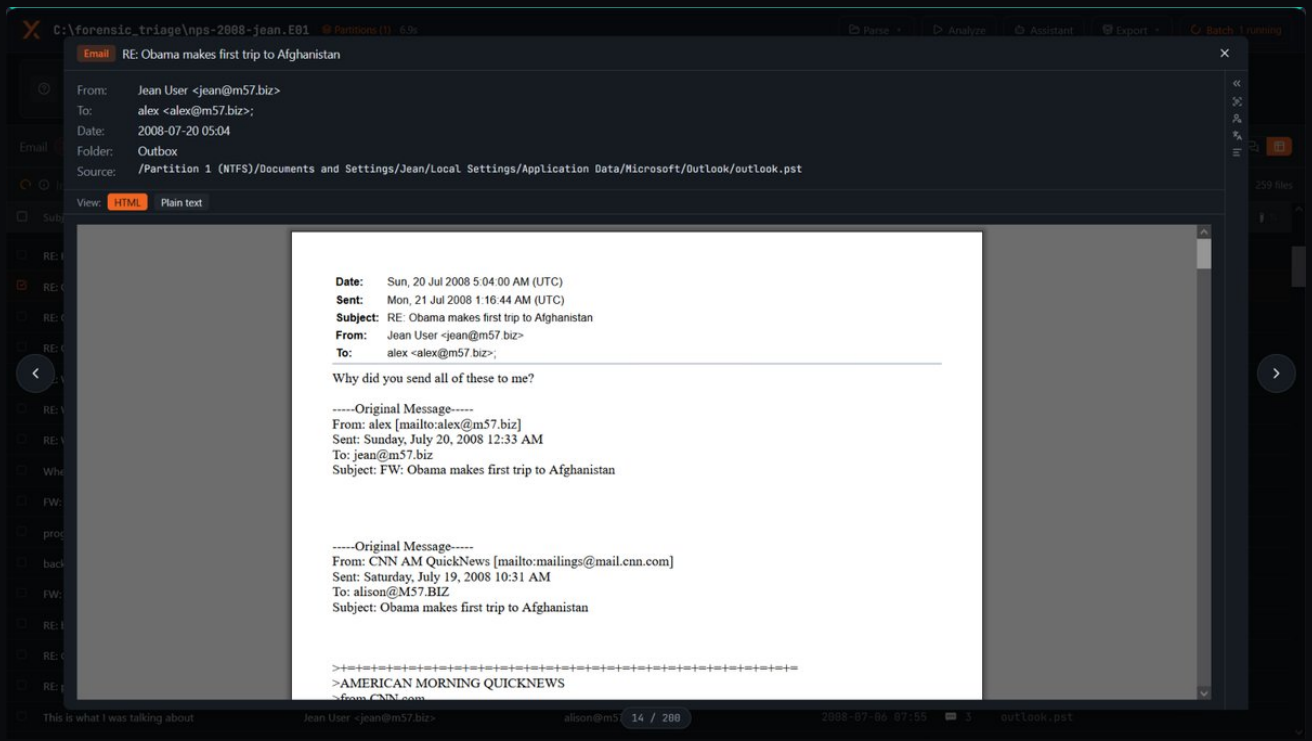

Click any row to open the email in a full-screen detail view showing its headers (From, To, Date, Folder, Source) and rendered body. Toggle between HTML and Plain text rendering, and use the arrows on either side of the overlay, or the position indicator at the bottom (e.g. 14 / 200), to step through the current result set without returning to the grid.

Searching Emails

Use the search bar at the top of the thread list to search by subject line or sender name across all parsed email sources.

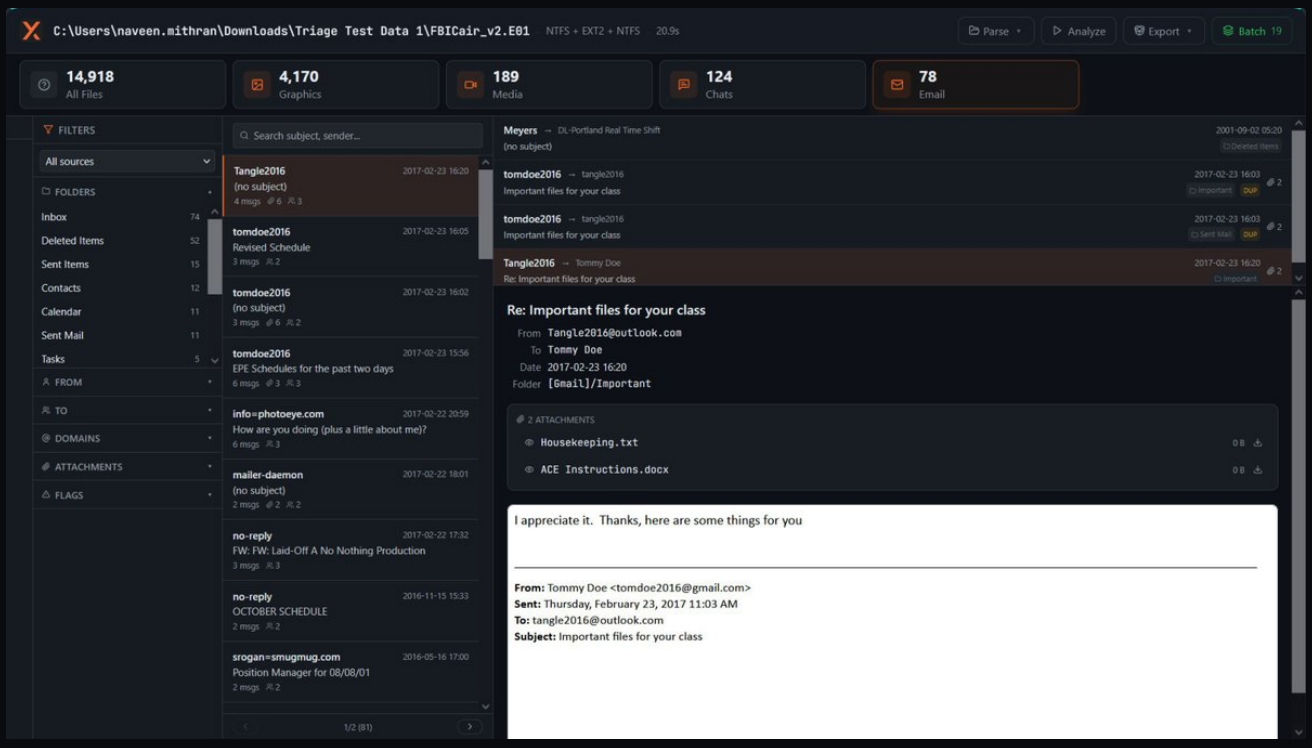

Results update instantly as you type. The search bar searches across all parsed PST and OST sources simultaneously, regardless of which source filter is active.

Email Filters

Available filters depend on the active view. Source is shared by both views from the top toolbar; Conversation view adds a left-hand filter panel, while Grid view adds Conversations only, Participants, and Date filters to the top toolbar. Active filters appear as badges. Source

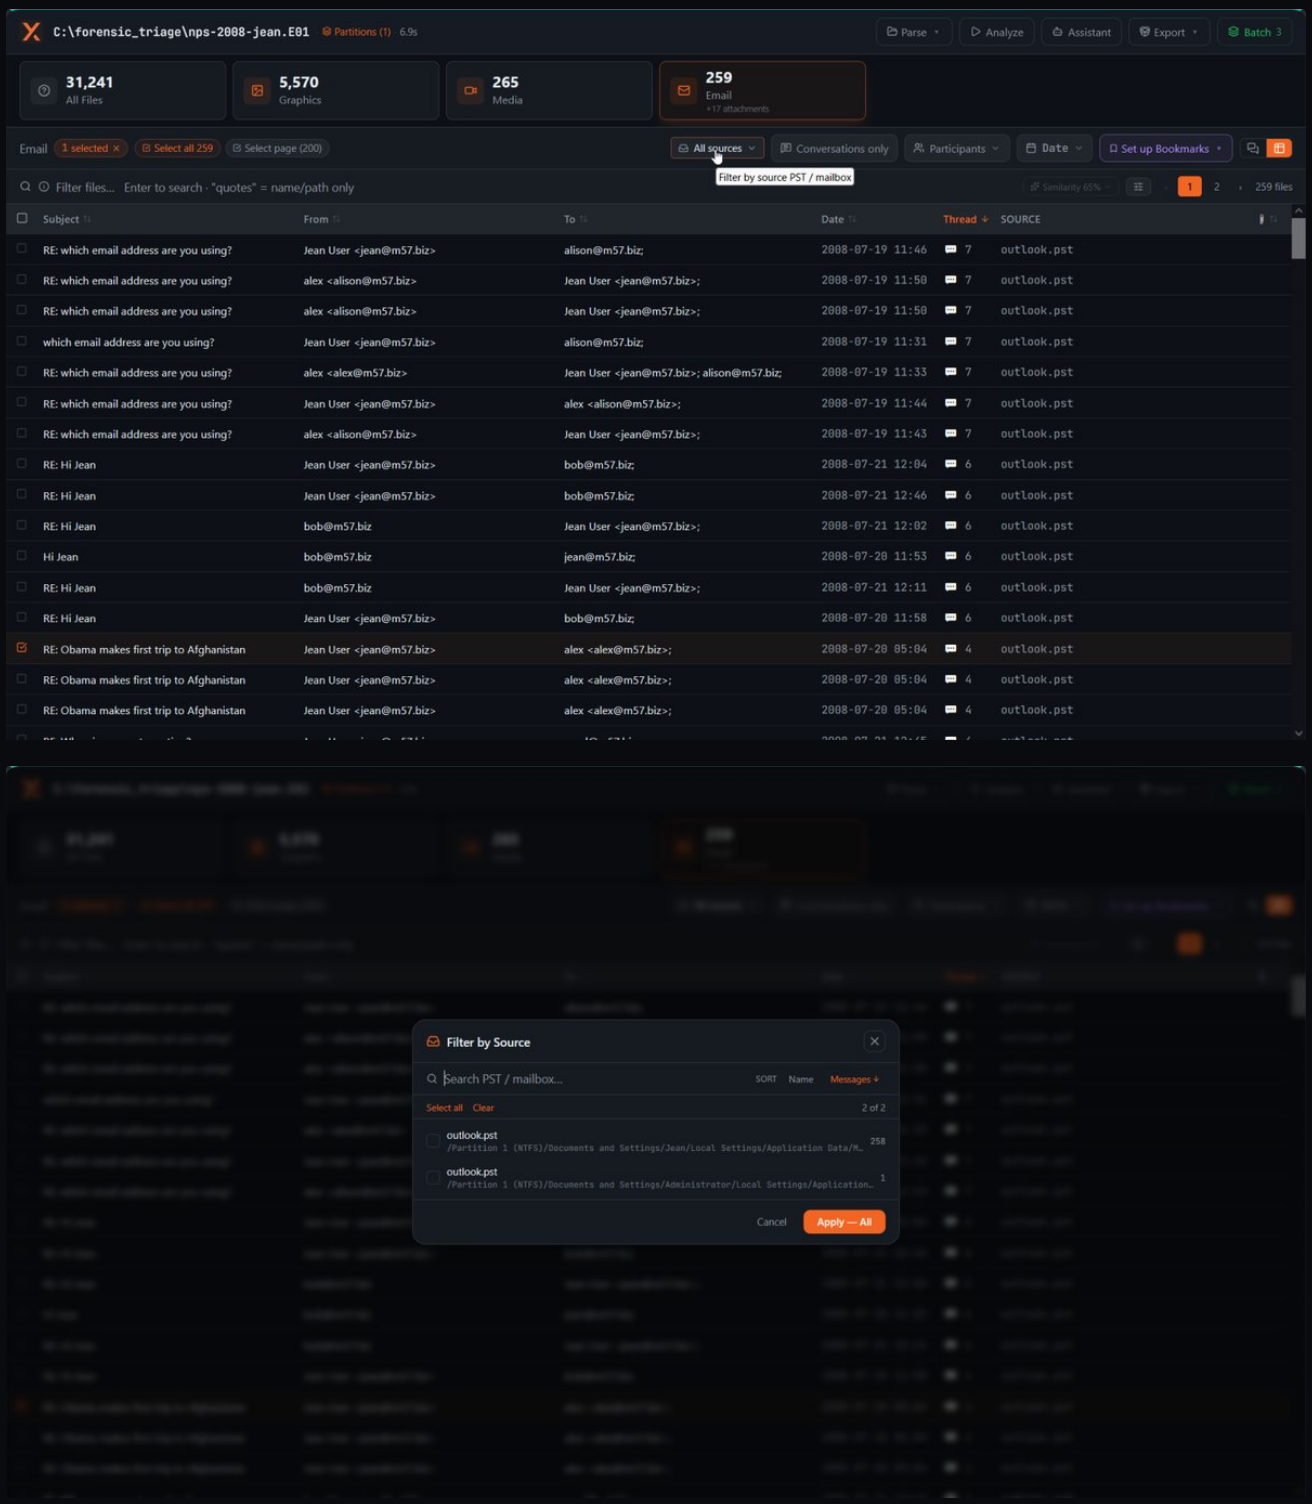

Search or sort the list of parsed PST and OST files, check the ones to include, and click Apply - All. Each source shows its message count. Leave all sources checked, or use Select all / Clear, to reset the filter.

From / To / Domains / Attachments (Conversation view) In Conversation view, the left panel lists From and To senders/recipients ranked by message count, Domains grouping emails by sending domain, and Attachments to isolate emails

with or without attachments. Click any entry to apply that filter to the thread list; expand a section's arrow to see its full list.

Conversations only (Grid view)

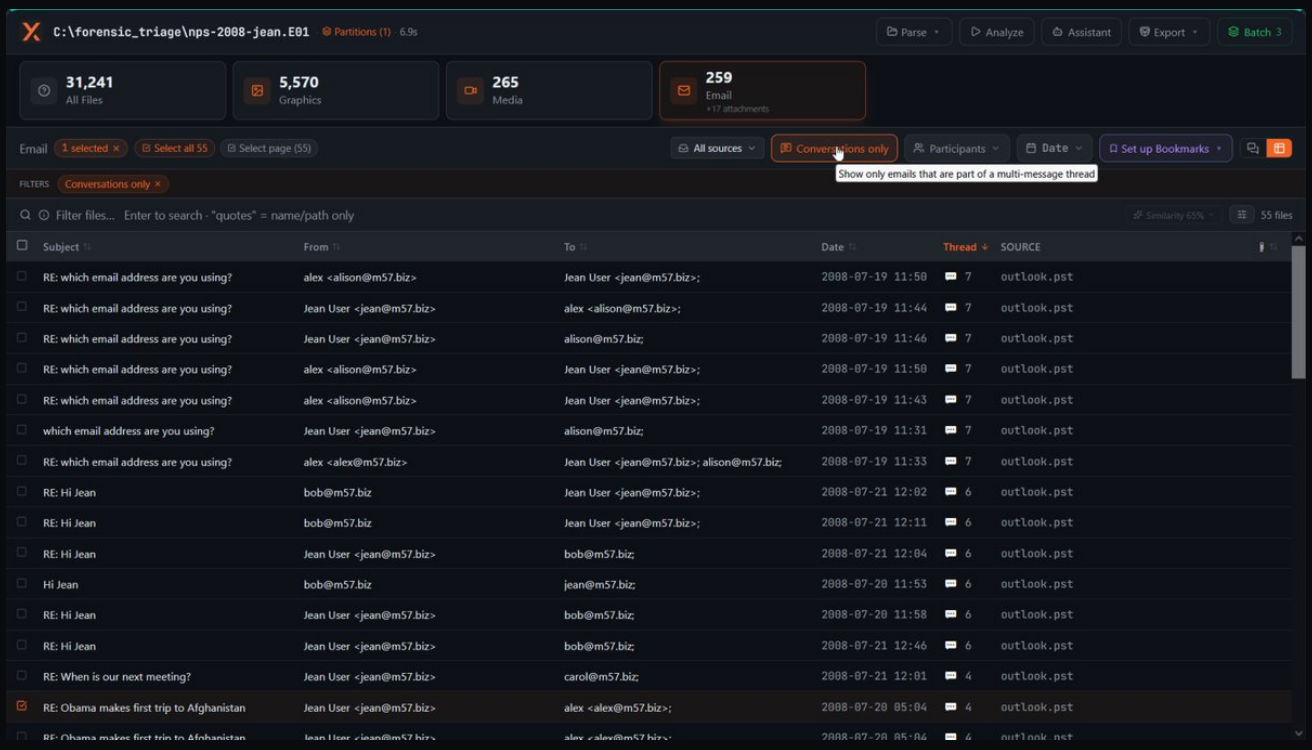

Since Grid view lists every email as an individual row, use Conversations only to filter out standalone messages and show only emails belonging to a thread of two or more. Participants (Grid view)

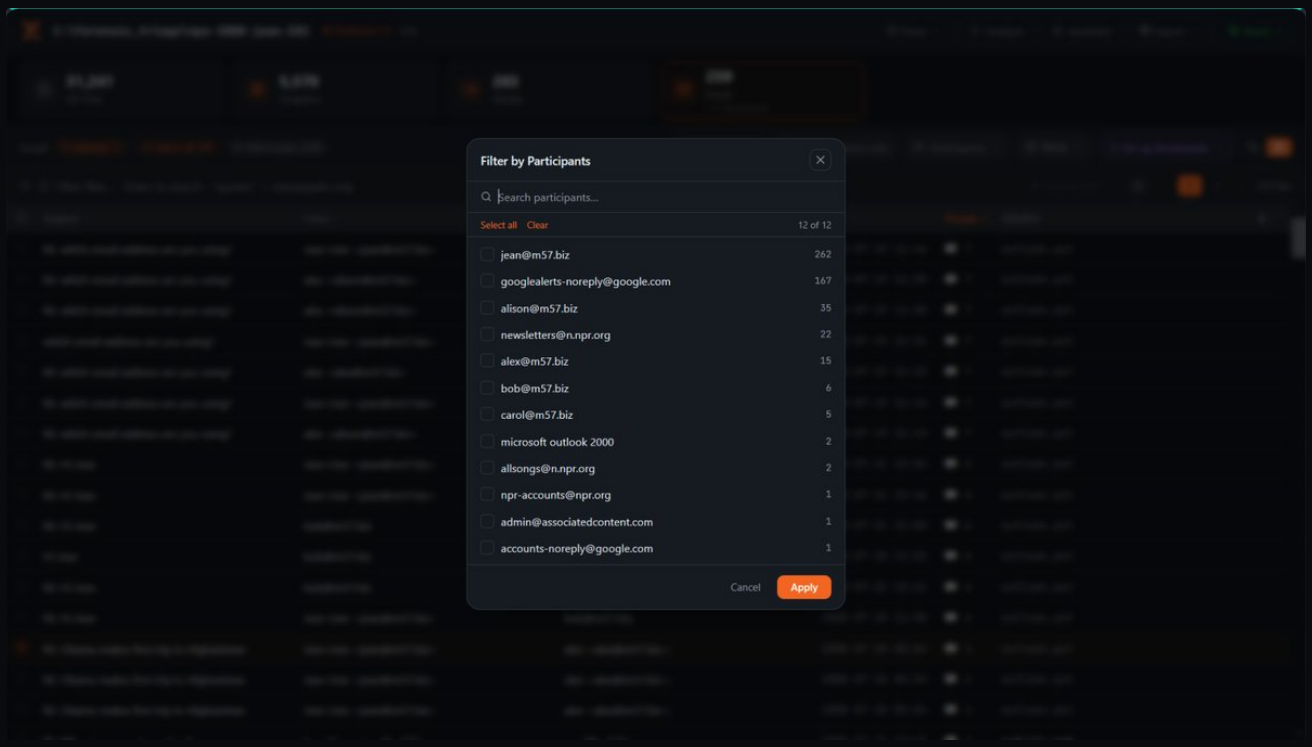

Click Participants to search or check any combination of senders and recipients across the case, each shown with a message count, then click Apply. This combines the From and To filters into a single multi-select search, useful when you need messages involving several specific people at once. Date (Grid view)

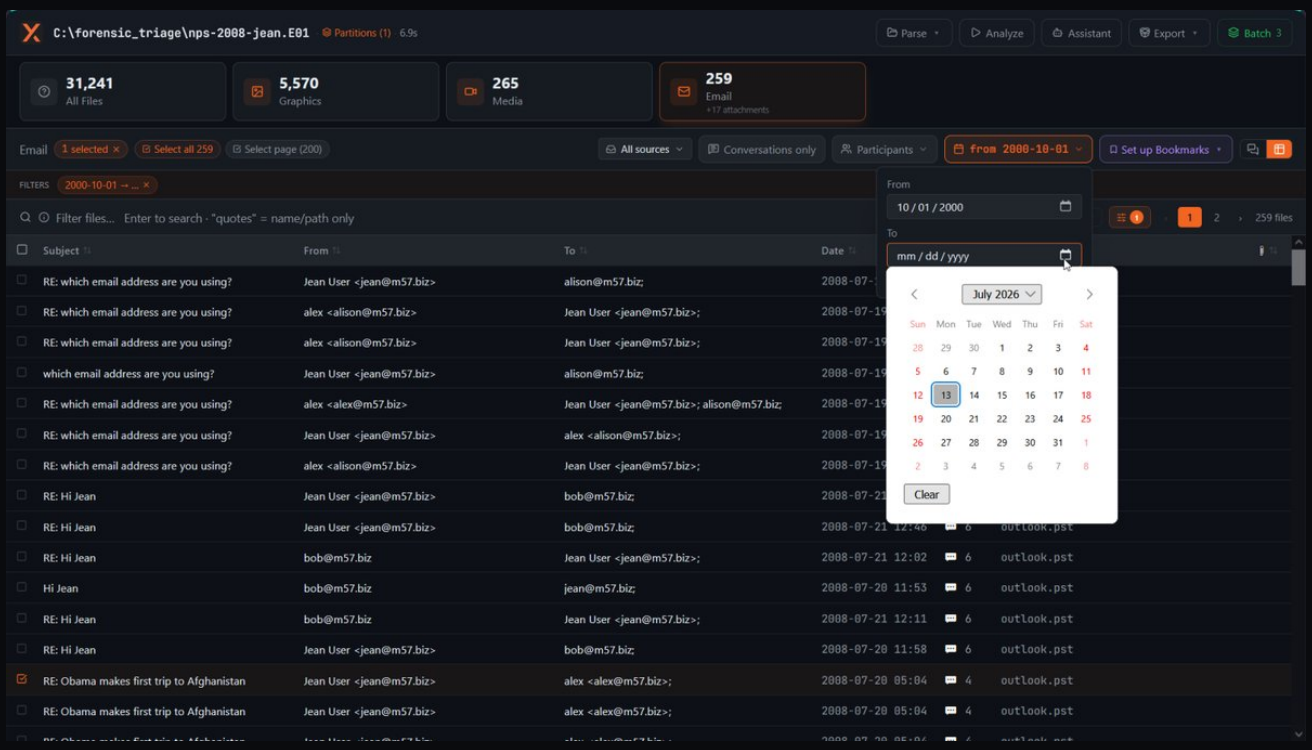

Click Date and set a From date, a To date, or both, either by typing or using the calendar. Click Clear to remove the date filter.

> Active filters from any of these controls appear as removable badges above the list. Combine multiple filters, e.g. a source plus a date range, to narrow results quickly.

PART 8

CHATS REVIEW

Parsing, reading, searching, and filtering extracted chat and messaging data

The Chats Review section lets investigators read and search all messaging data extracted from the evidence. It supports WhatsApp, SMS/MMS, Facebook Messenger, and other chat applications. Why this matters: Chat data from mobile devices can contain hundreds of thousands of messages across multiple apps.

ARMOURop consolidates all messaging sources into one interface with conversation-level and in-conversation search.

How To Access

Click the Chats tab in the file category summary bar. Chat data must be parsed first.

Parsing chat databases

Before reviewing chat data, parse the messaging databases found in the evidence. Click Parse in the top-right toolbar and select Parse Chats.

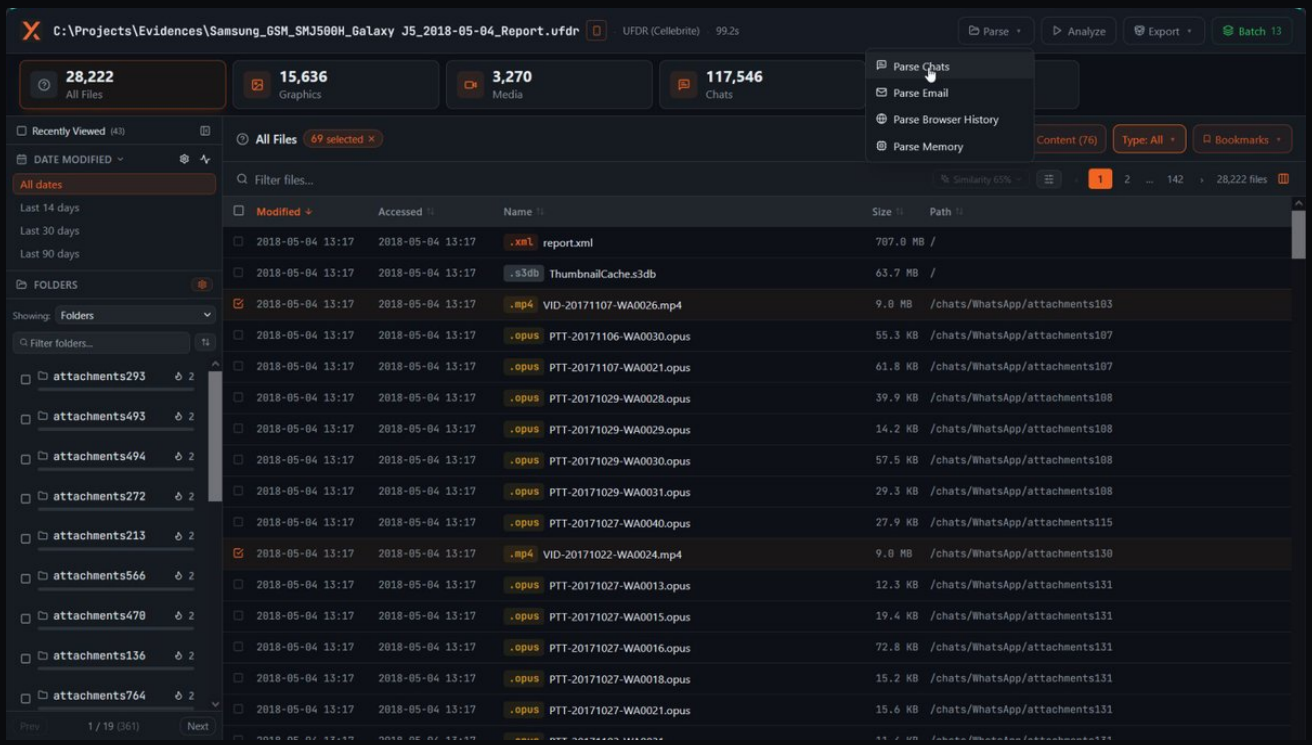

The Parse Chats dialog lists all messaging app databases discovered in the evidence. Each entry shows the app name, database filename, and file path. Select the databases you want to parse and click Parse Databases.

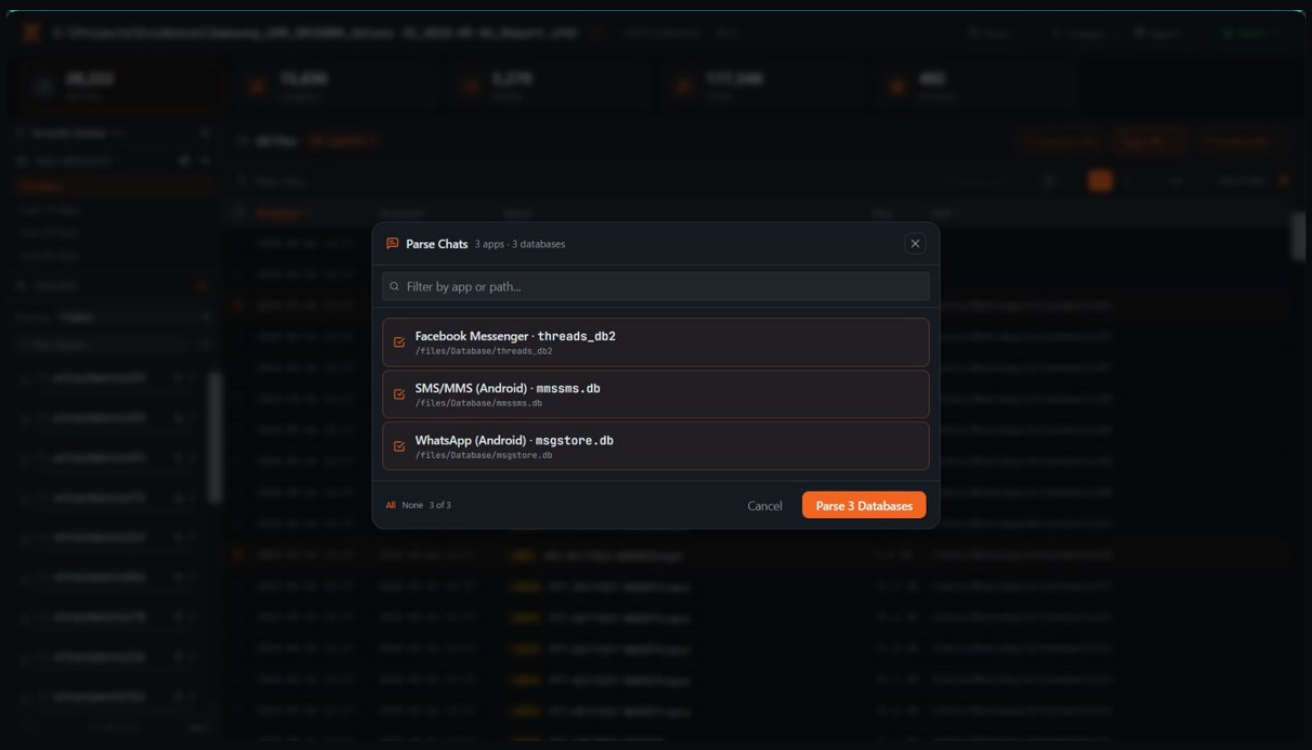

ARMOURop automatically identifies the app associated with each database file. Supported apps include WhatsApp, Facebook Messenger, SMS/MMS (Android), and others found in mobile extractions.

Reviewing conversations

Once parsed, the Chats tab offers two view modes: Conversation view, a three-panel layout for reading message threads in context, and Grid (table) view, a spreadsheet-style list for scanning and batch-selecting large volumes of conversations. Switch between them using the two icon buttons at the top right of the Chats tab.

Conversation View

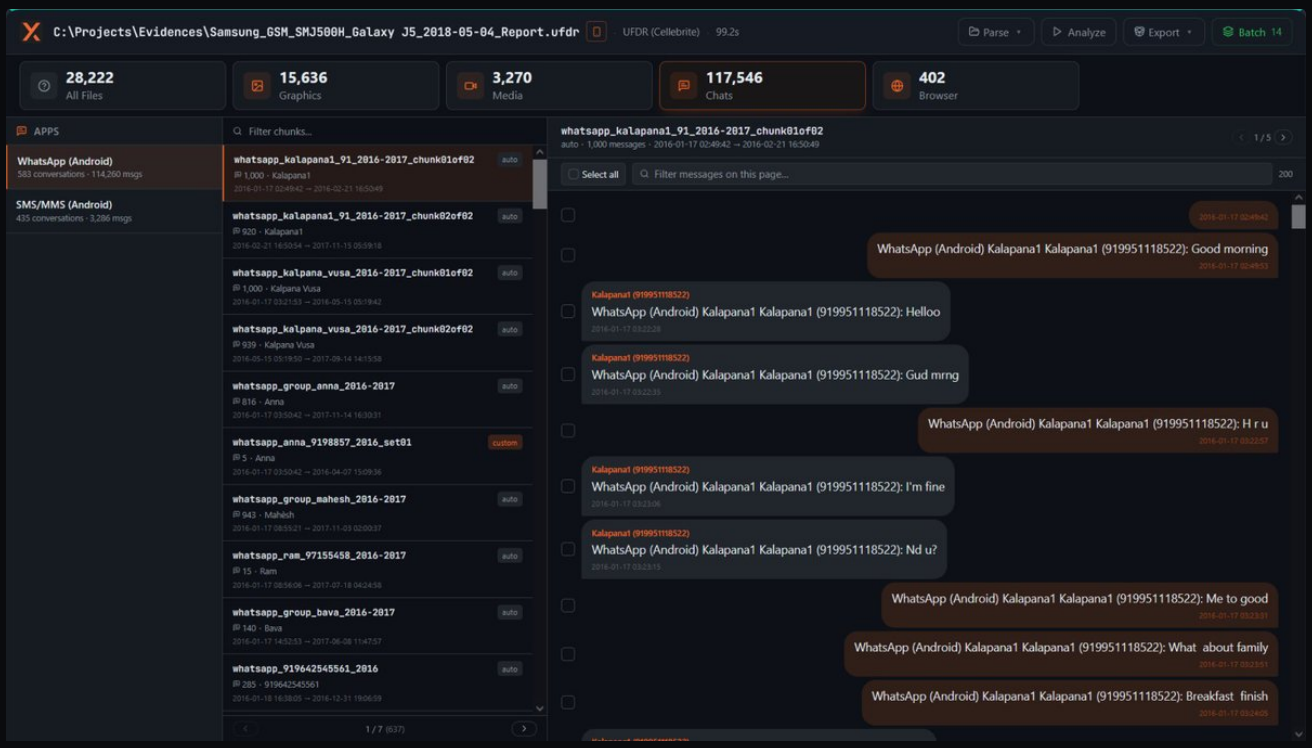

Apps panel All parsed messaging apps are listed with total conversation and message counts. Click an app to view only its conversations in the middle panel. Conversation chunks Large conversations are split into chunks for performance. Each chunk shows the contact name, message count, and the date range it covers. Click any chunk to open it in the message reader on the right.

Grid (Table) View

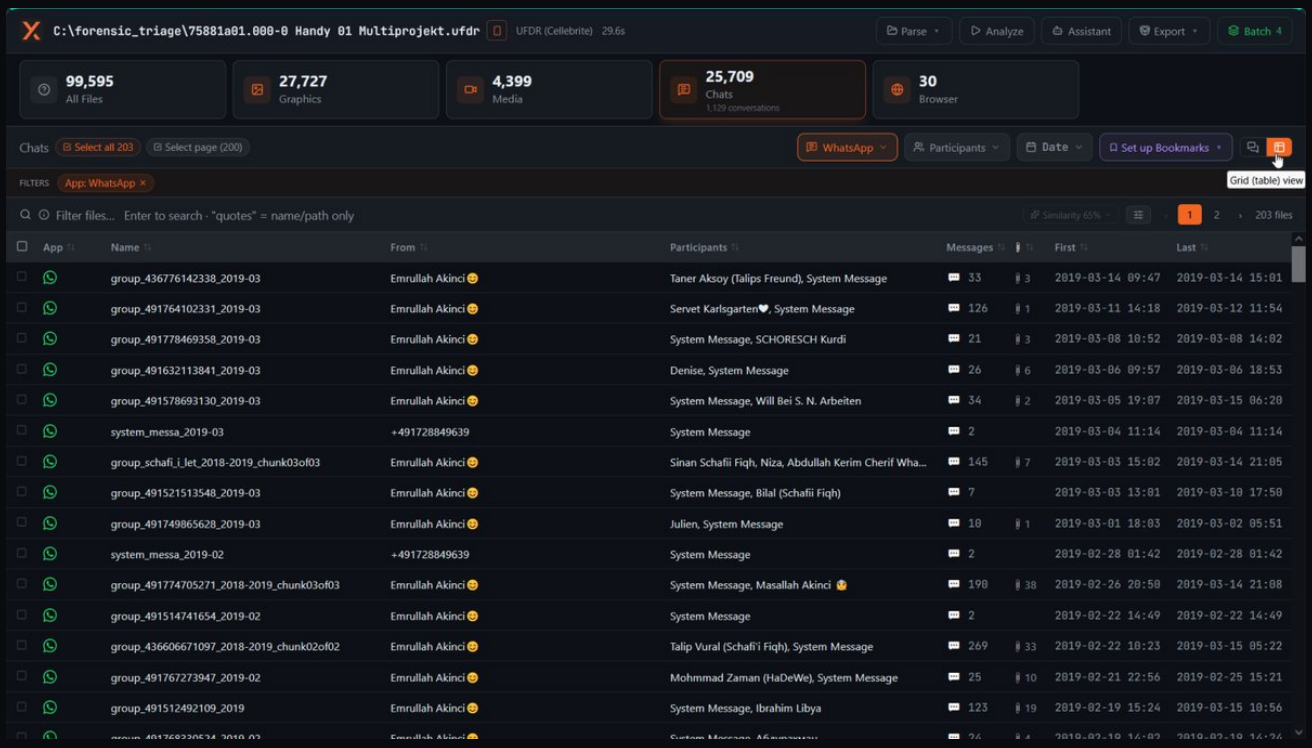

Grid view lists every conversation chunk individually rather than grouped by app, making it well suited to sorting, bulk selection, and export. Click any column header to sort by that column. Use the checkboxes to select individual chunks, or Select all / Select page above the list to select in bulk, then use Set up Bookmarks to tag the selection.

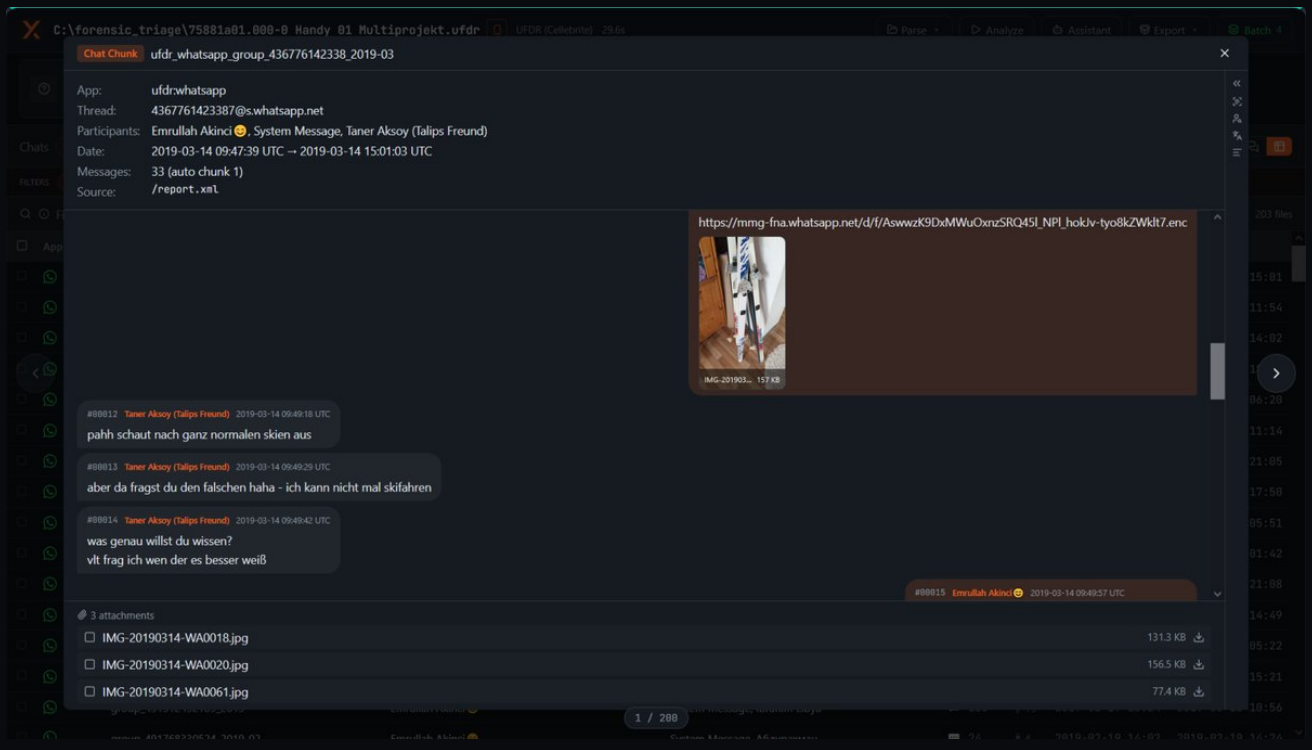

Click any row to open the chunk in a full-screen detail view showing its headers (App, Thread, Participants, Date range, Messages, Source) and the rendered message bubbles, including inline image attachments. A list of attachments with file sizes and download links appears below the conversation.

Use the arrows on either side of the overlay, or the position indicator at the bottom (e.g. 1 / 200), to step through the current result set without returning to the grid.

Searching Conversations

Two levels of search are available: searching across conversations to find a contact, and searching within a conversation to find a specific message. Search conversations by contact

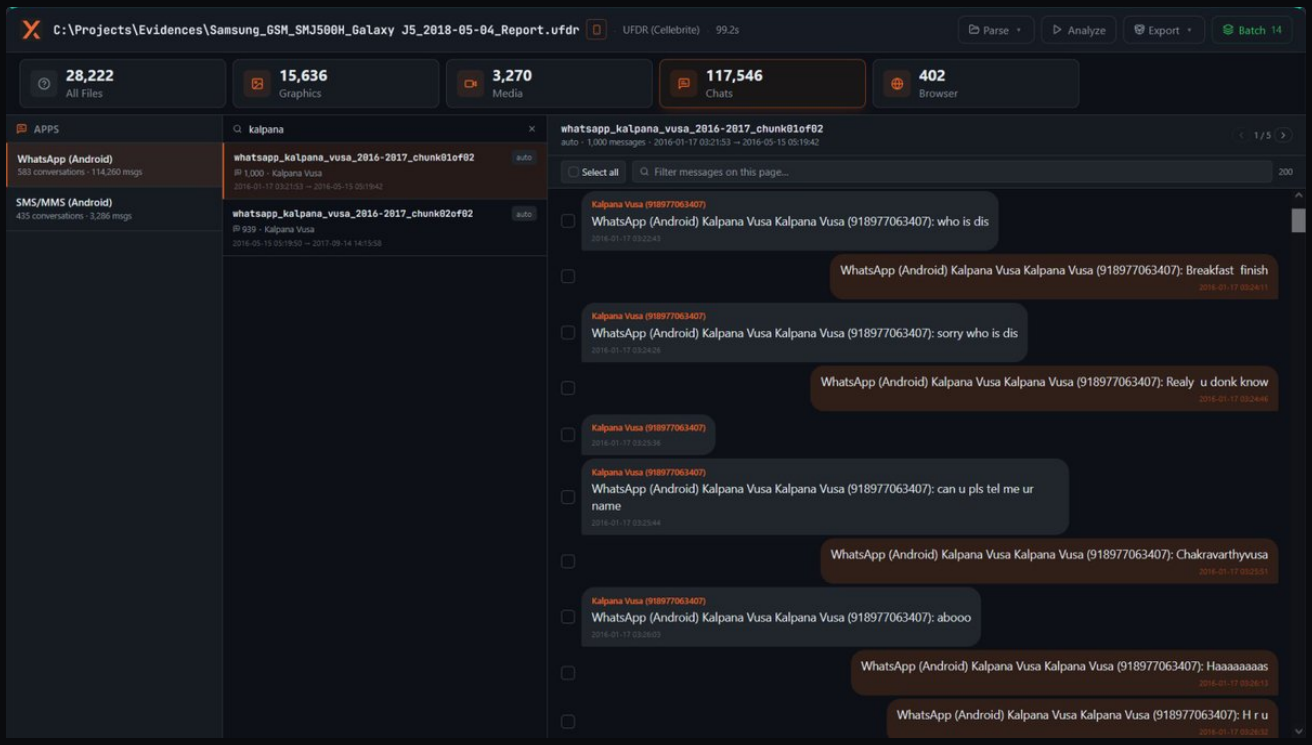

Use the search bar above the conversation list to filter chunks by contact name. Only conversations involving that contact are shown. Search within a conversation

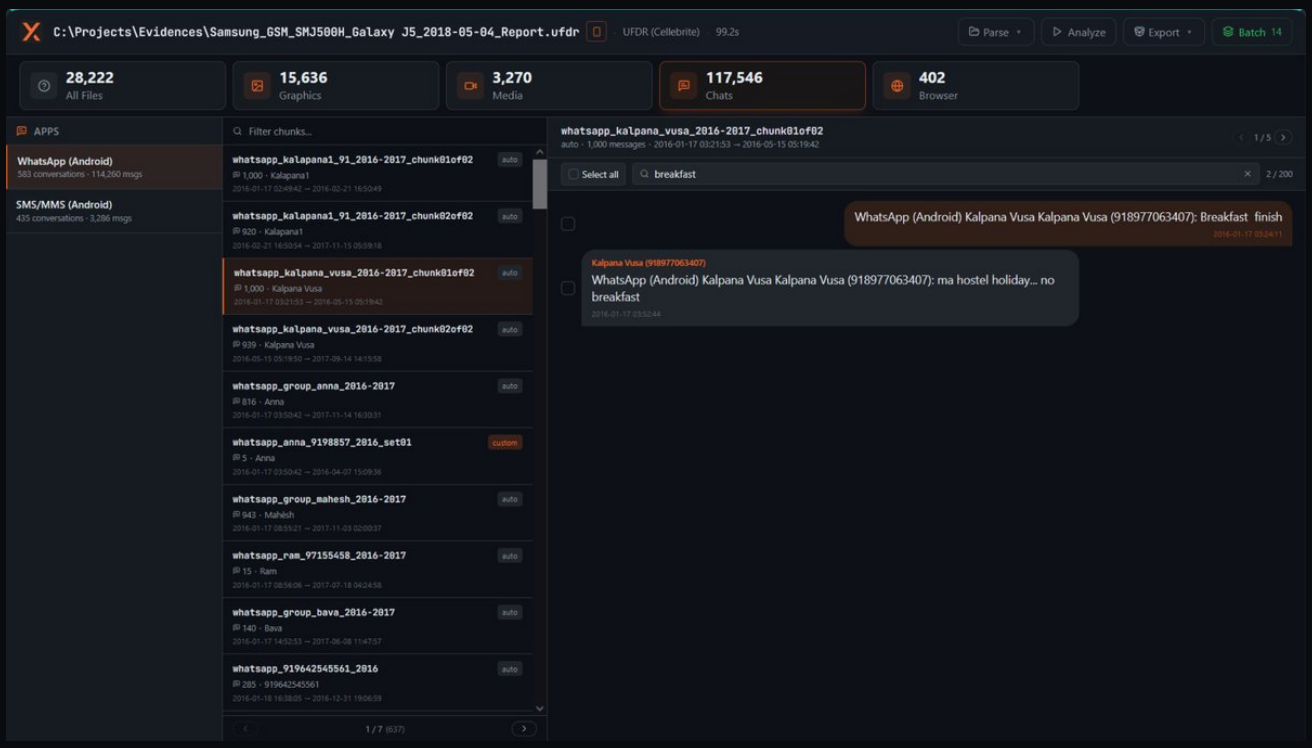

Use the message filter bar on the right panel to search for a specific word or phrase within the selected chunk. Only messages containing the search term are shown, with the result count displayed (e.g. 2 of 200 messages).

Use both search levels together for maximum precision. First narrow to the contact, then search for the keyword inside their messages.

Chat Filters

Grid (table) view adds three filters to the top toolbar: App, Participants, and Date. Active filters appear as badges above the list. App

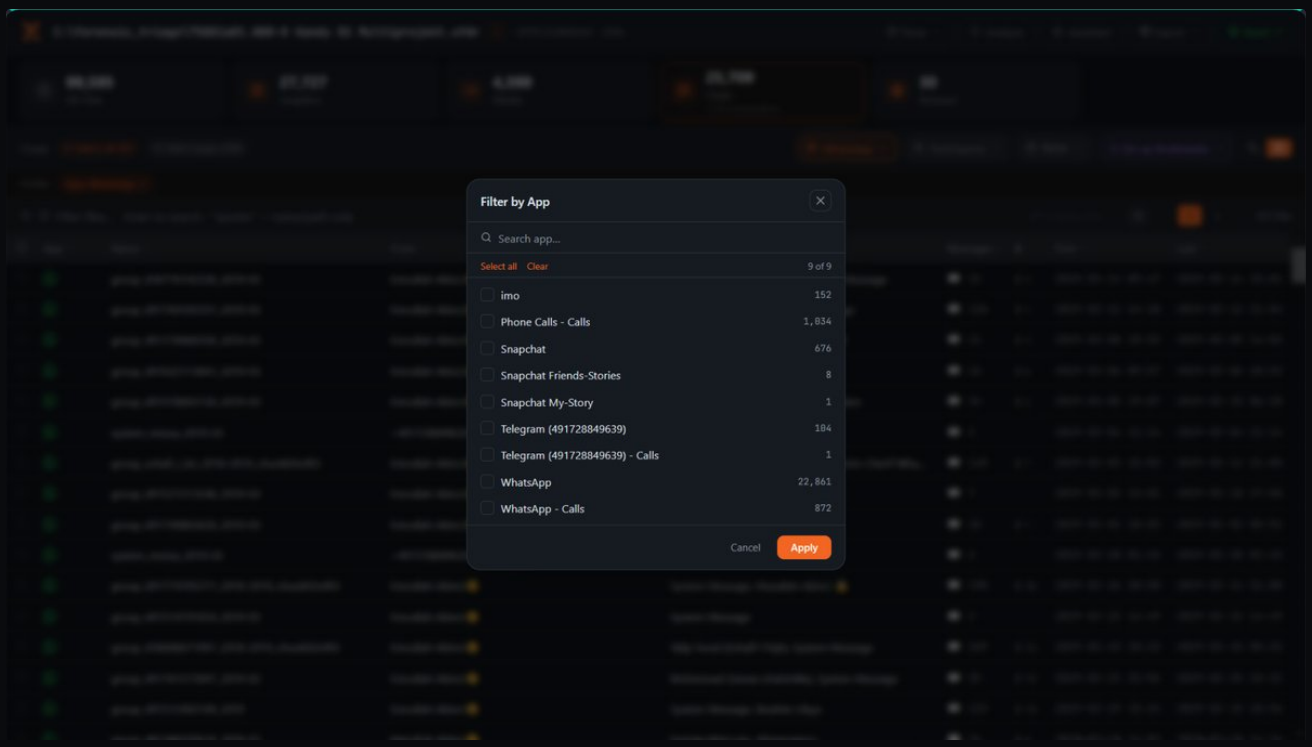

Search or check the apps to include (e.g. WhatsApp, Telegram, Snapchat, Phone Calls), then click Apply. Use Select all / Clear to reset the filter. This replaces manually clicking through the apps panel used in Conversation view. Participants

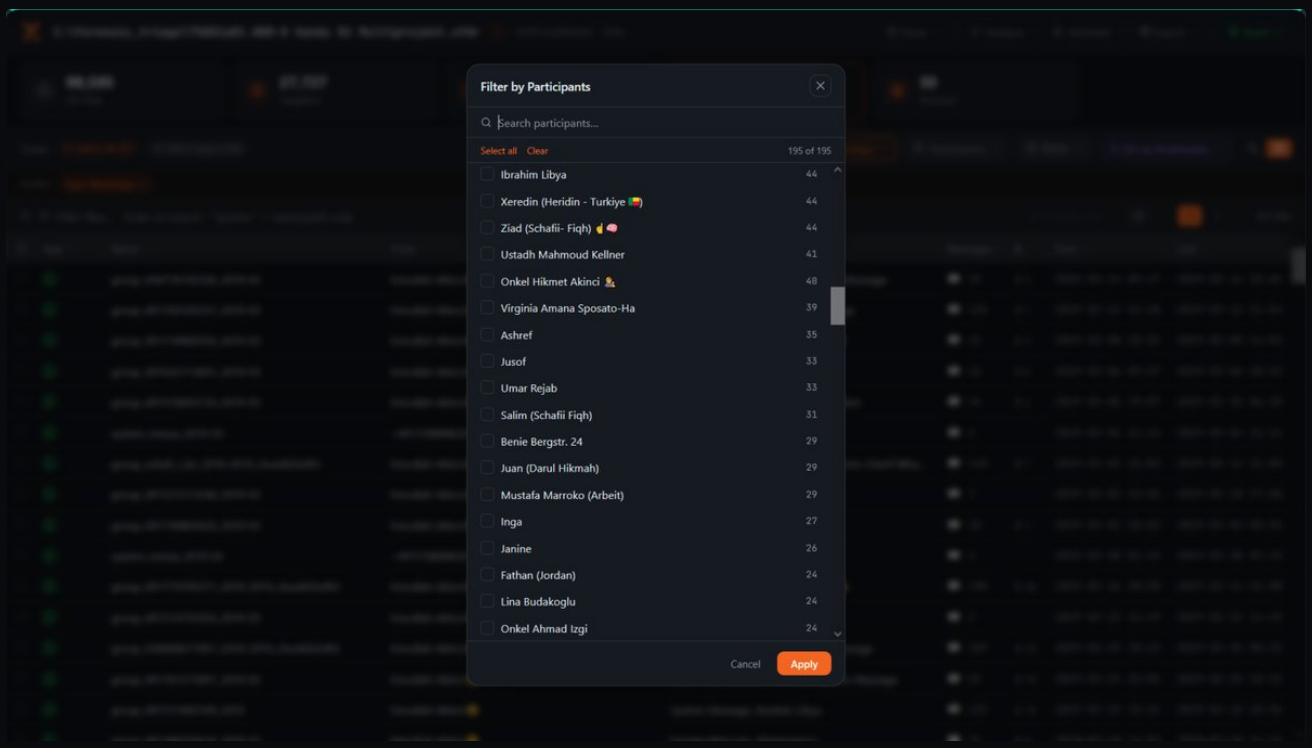

Click Participants to search or check any combination of contacts across the case, each shown with a message count, then click Apply. Useful for surfacing every conversation chunk involving one or more specific people, regardless of which app or group chat they appear in. Date

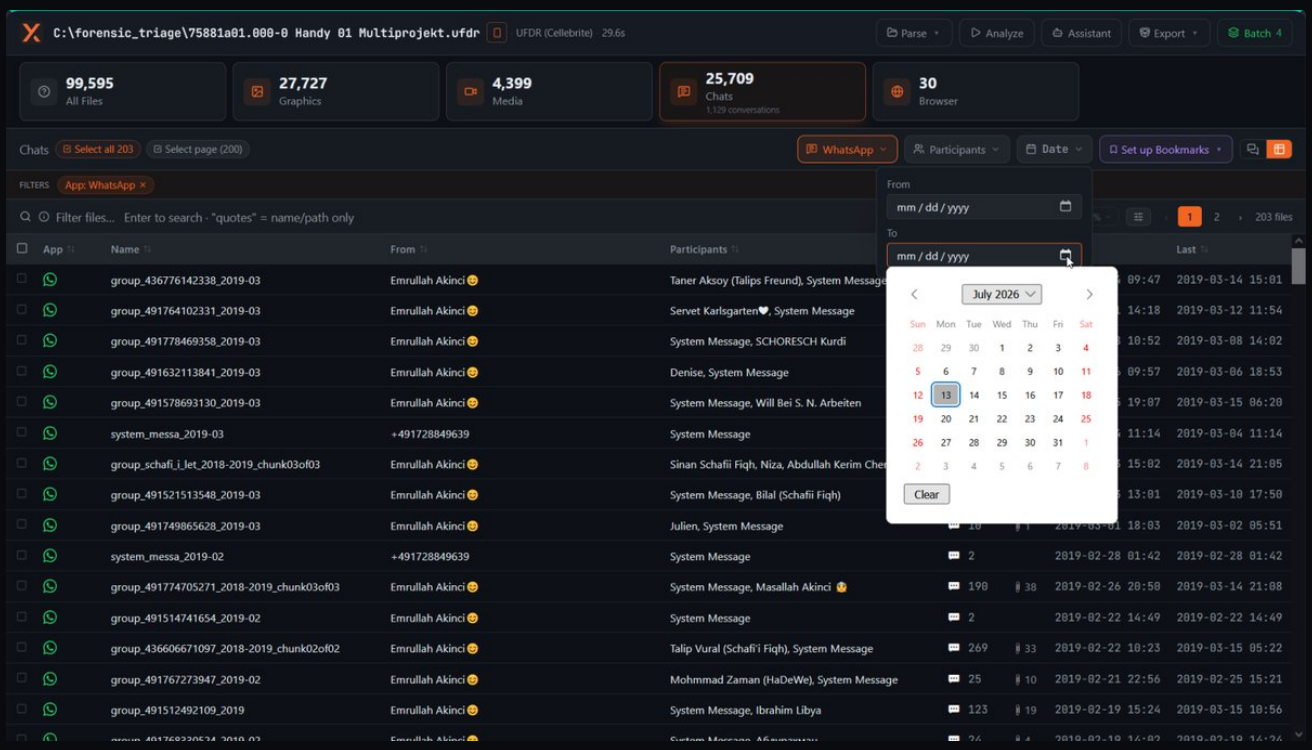

Click Date and set a From date, a To date, or both, either by typing or using the calendar. Click Clear to remove the date filter.

> Combine App, Participants, and Date together to narrow a large multi-app extraction down to a specific conversation window quickly, e.g. all WhatsApp chunks involving a given contact within a date range.

PART 9

BROWSER REVIEW

Parsing and reviewing browser history, downloads, and search activity

The Browser Review section presents all web browsing history extracted from the evidence, organized by AI-assigned category. Why this matters: Browsing history can reveal intent, research patterns, and financial activity. Categories like Banking/Finance, Streaming, and URL Shorteners let investigators zero in on suspicious activity in seconds.

How To Access

Click the Browser tab in the file category summary bar. Browser history must be parsed first.

Parsing browser history

Before reviewing browser history, parse the browser databases found in the evidence. Click Parse in the top-right toolbar and select Parse Browser History.

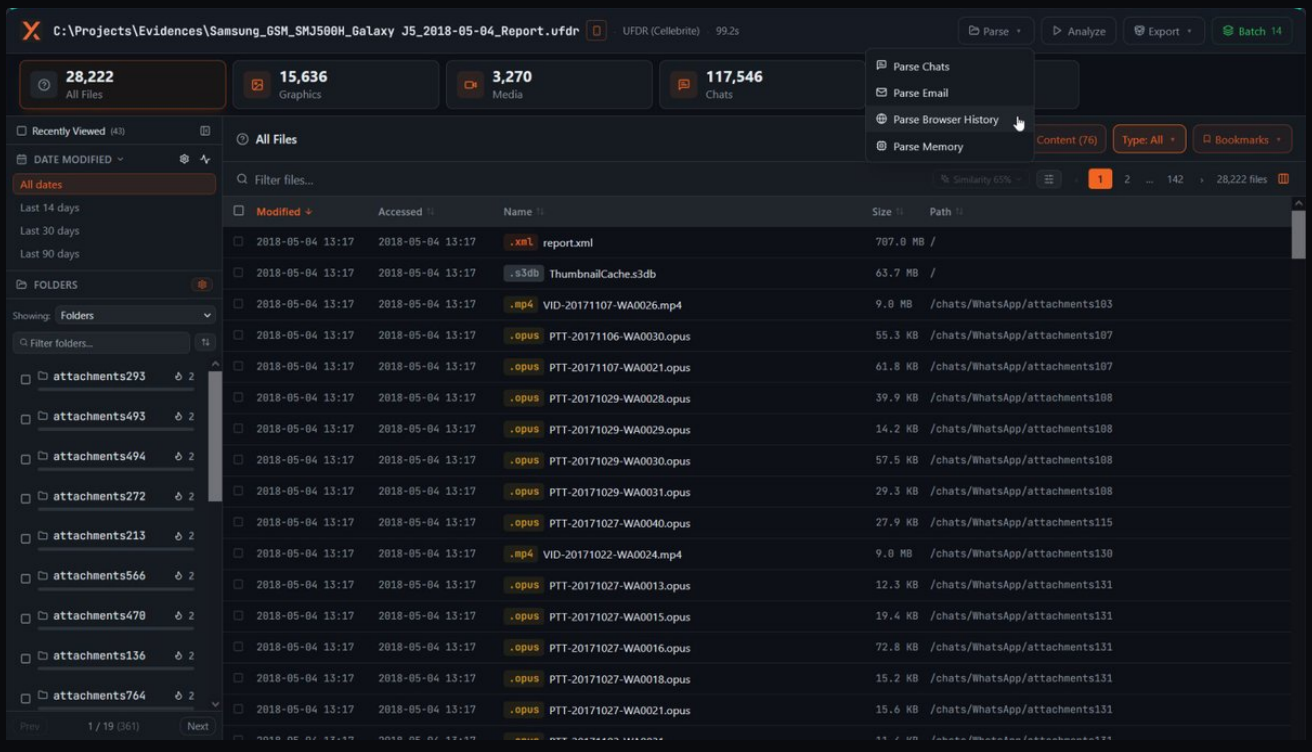

The dialog lists all browsers detected in the evidence, grouped by browser type with a database count for each. Select the browsers to parse and click Parse Browsers.

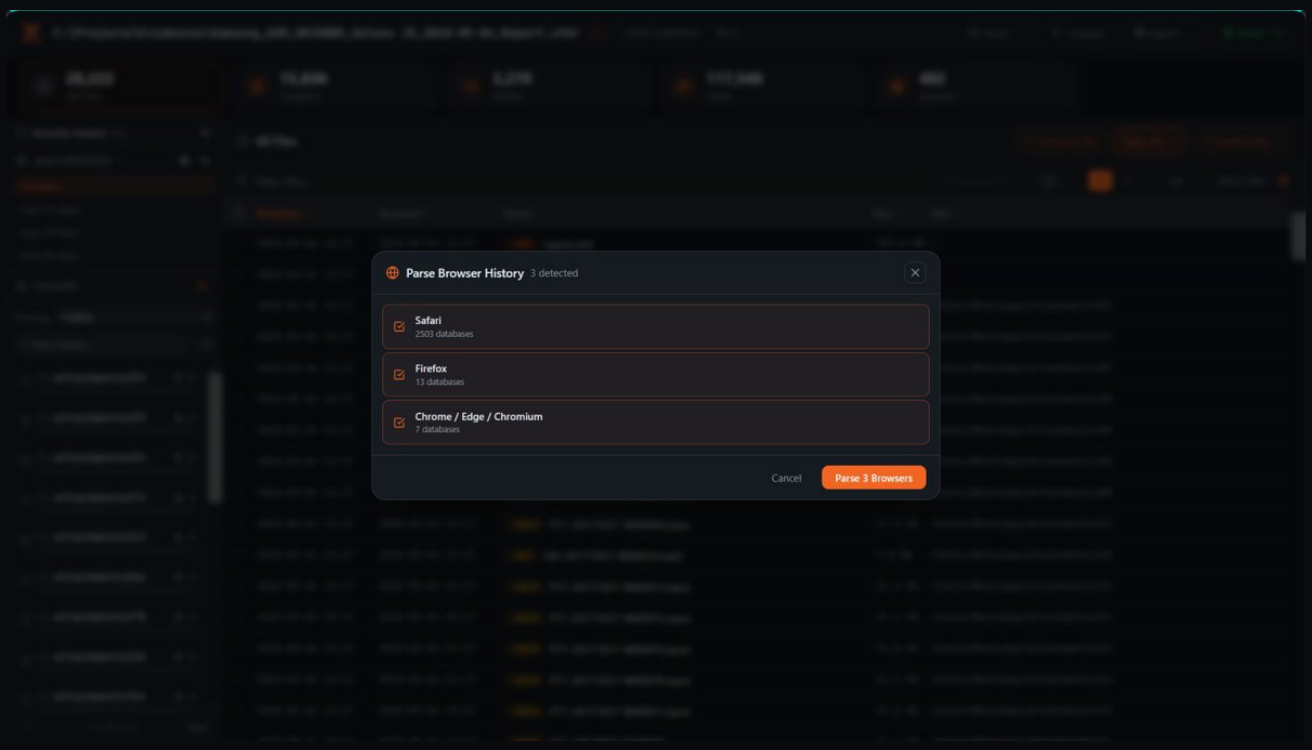

ARMOURop detects and parses history from Safari, Firefox, Chrome, Edge, and Chromium-based browsers.

Reviewing history

Once parsed, the Browser tab displays all history entries in a two-panel layout.

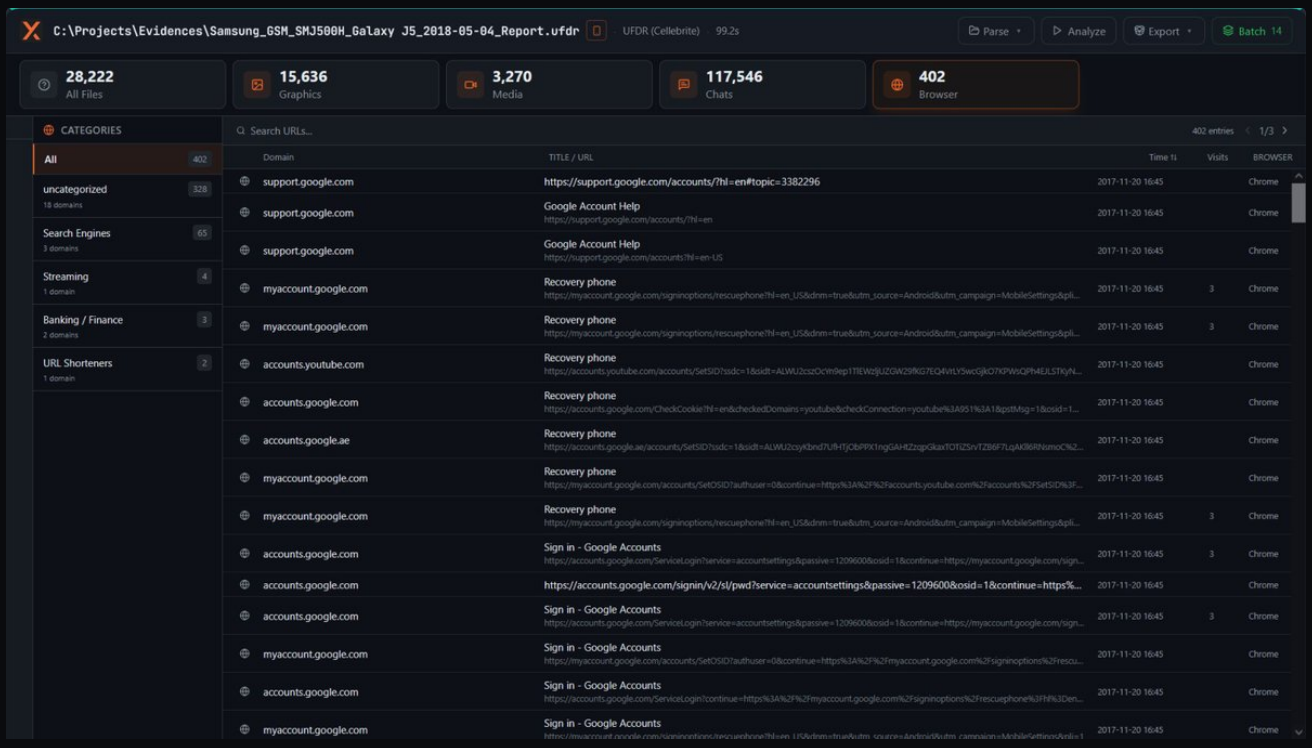

The left panel shows AI-assigned website categories with entry counts. The right panel shows the full history list with domain, page title, URL, timestamp, visit count, and the browser it came from.

Categories

The AI automatically categorizes every URL in the browsing history. Click any category to instantly filter the history list to only entries from that type of website.

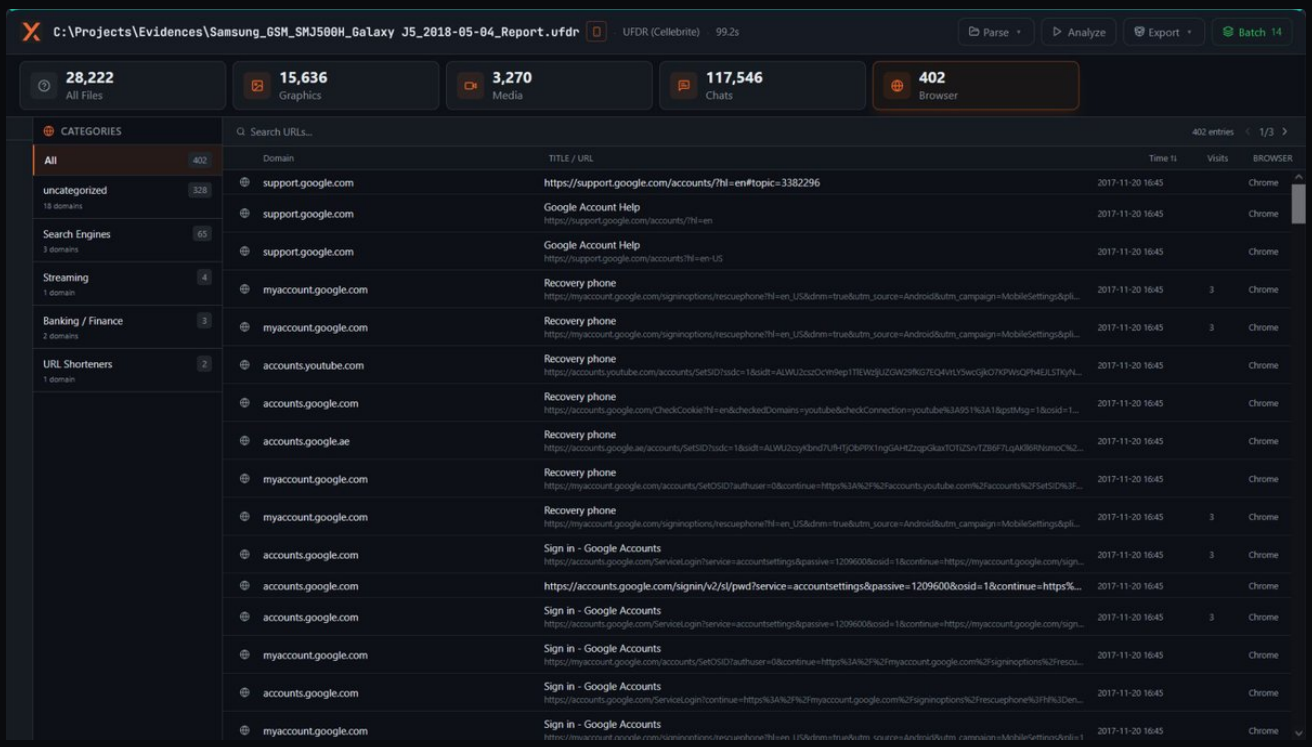

All Shows the complete browsing history across all categories and browsers. Uncategorized URLs the AI could not assign to a known category. Often the largest group, worth reviewing for unusual domains. Search Engines Searches performed across Google, Bing, and other platforms. Useful for establishing what the subject was researching.

Streaming Video and audio streaming platforms such as YouTube and Netflix. Banking / Finance Financial services websites including banks, payment platforms, and money transfer services. High-value for fraud investigations. URL Shorteners Visits to link-shortening services. Worth investigating as they can obscure the true destination of a link.

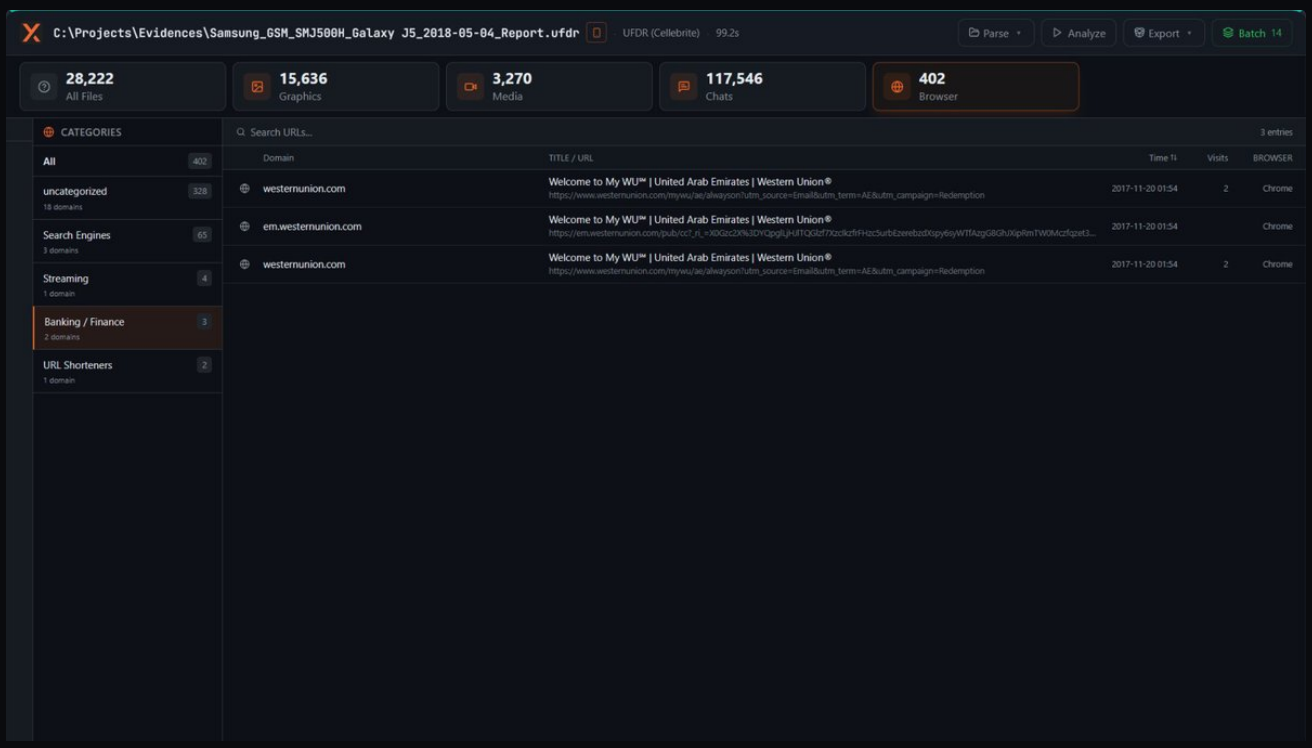

Searching Browser History

Use the search bar at the top of the history list to search across all URL and page title text.

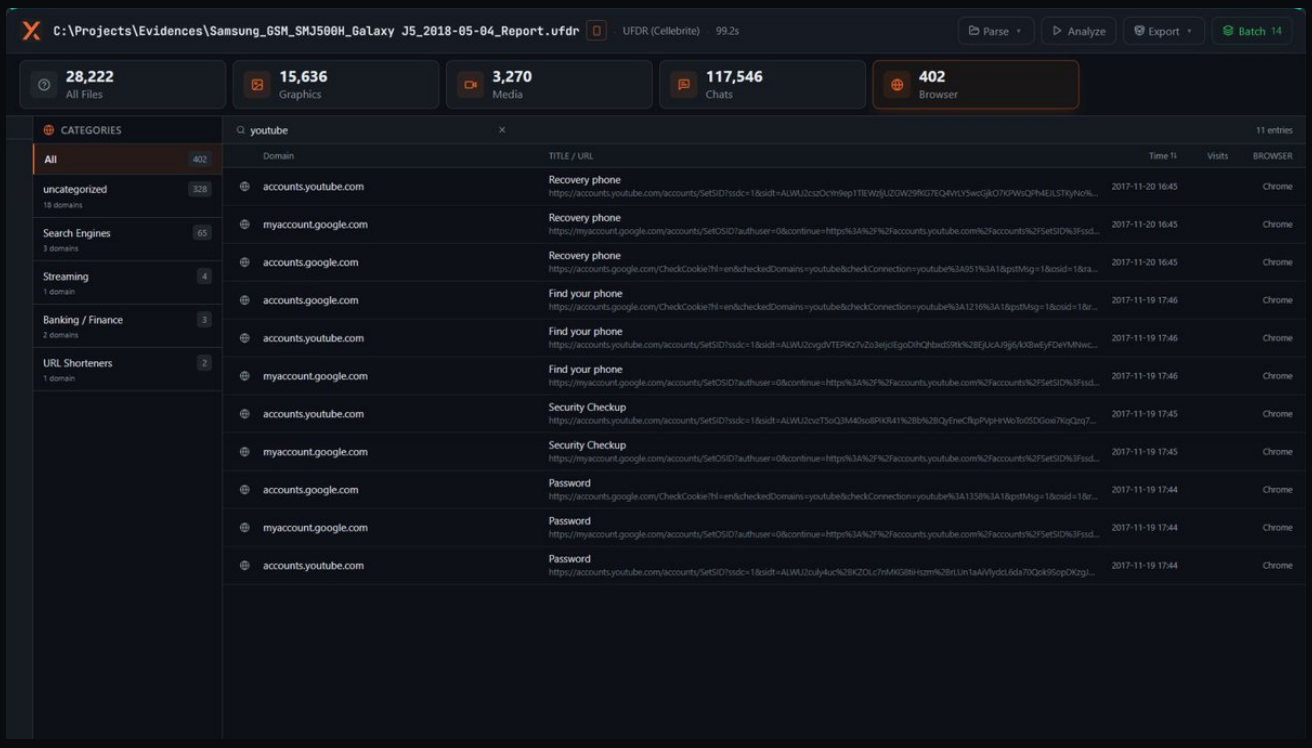

Results update instantly as you type, showing only entries that contain the search term anywhere in the domain, title, or full URL.

> Use categories and URL search together. For example, filter to Search Engines then search for a keyword to see exactly what the subject was looking up on that topic.

PART 10

AI ASSISTANT

Interrogating a case conversationally using natural-language queries

A case chat assistant that can query, search, bookmark, export, and run batches against the current case using natural language, powered by a local Ollama model configured in Settings → AI Assistant.

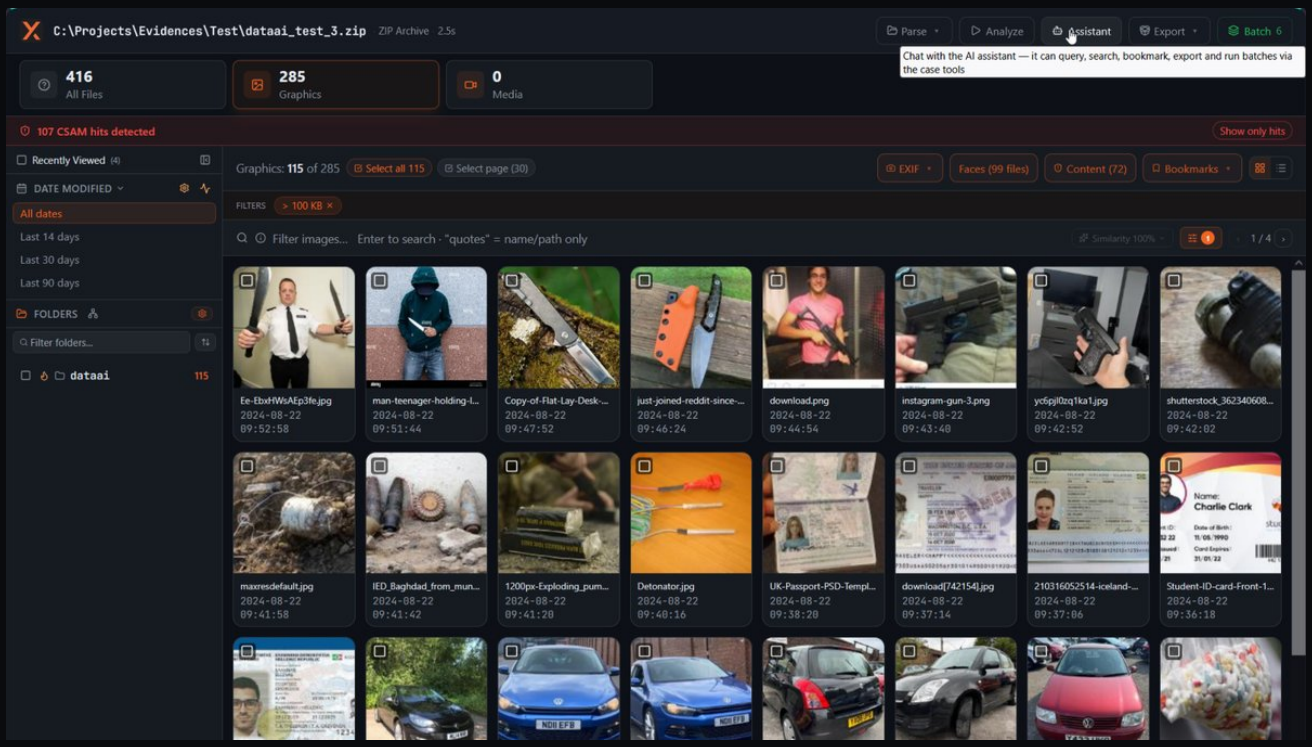

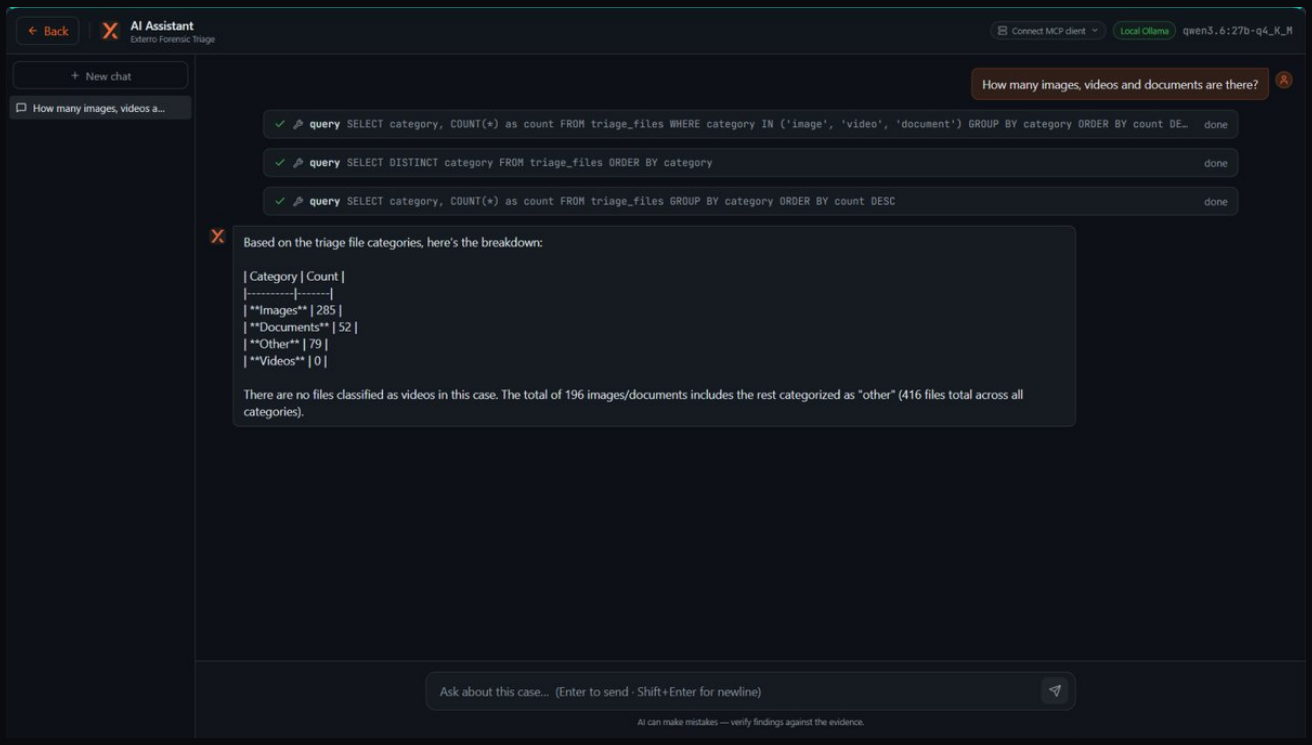

1. Open the AI Assistant Click Assistant in the Triage Console toolbar to open the chat interface for the current case. Use Back at any time to return to the Triage Console.

2. Ask a question in natural language Type a question about the case into the input field at the bottom, e.g. "How many images, videos and documents are there?" Press Enter to send, or Shift+Enter for a new line.

3. Watch the assistant work The assistant calls case tools such as query as needed to answer, shown inline with the query text and a done status once each call completes. Several tool calls can chain together to answer one question.

4. Review the answer The assistant summarizes tool results in a readable response, including tables where useful. Always verify important findings against the underlying evidence.

5. Manage chats Click New chat to start a fresh conversation, or select a past conversation from the sidebar to continue it. Connect MCP client Connects the assistant to an external MCP client for additional tool access beyond the built-in case tools. Local Ollama / model badge Shows the active connection type and the model currently in use (e.g. qwen3.6:27b-q4_K_M), as configured in Settings → AI Assistant.

> AI can make mistakes - verify findings against the evidence. The assistant requires a tool-calling capable model; if it stops using case tools mid-conversation, check the model and connection in Settings → AI Assistant. WHEN WOULD I USE THIS?
>
> When you want a fast, conversational way to interrogate a case, such as getting file-category breakdowns, searching for specific artifacts, or triggering bookmarks and exports, without manually building queries or navigating the Triage Console yourself.

PART 11

BOOKMARKS

Tagging, applying, and filtering files with colour-coded bookmark labels

Bookmarks let investigators tag files with color-coded labels during review. Each bookmark has a name, a color, and a keyboard hotkey, so files can be tagged instantly without interrupting the review flow. Why this matters: When reviewing thousands of files, investigators need a fast way to mark key evidence.

Bookmarks let you tag a file with a single keystroke, assign it to a category, and instantly recall all tagged files at any point. WHEN WOULD I USE THIS? During any review session where you want to flag files for follow-up, categorize evidence by relevance, or mark items to include in a report. Bookmarks persist for the life of the case.

Setting up bookmarks

Click the Set up Bookmarks button in the top-right toolbar, then select Manage Bookmarks to open the bookmark manager.

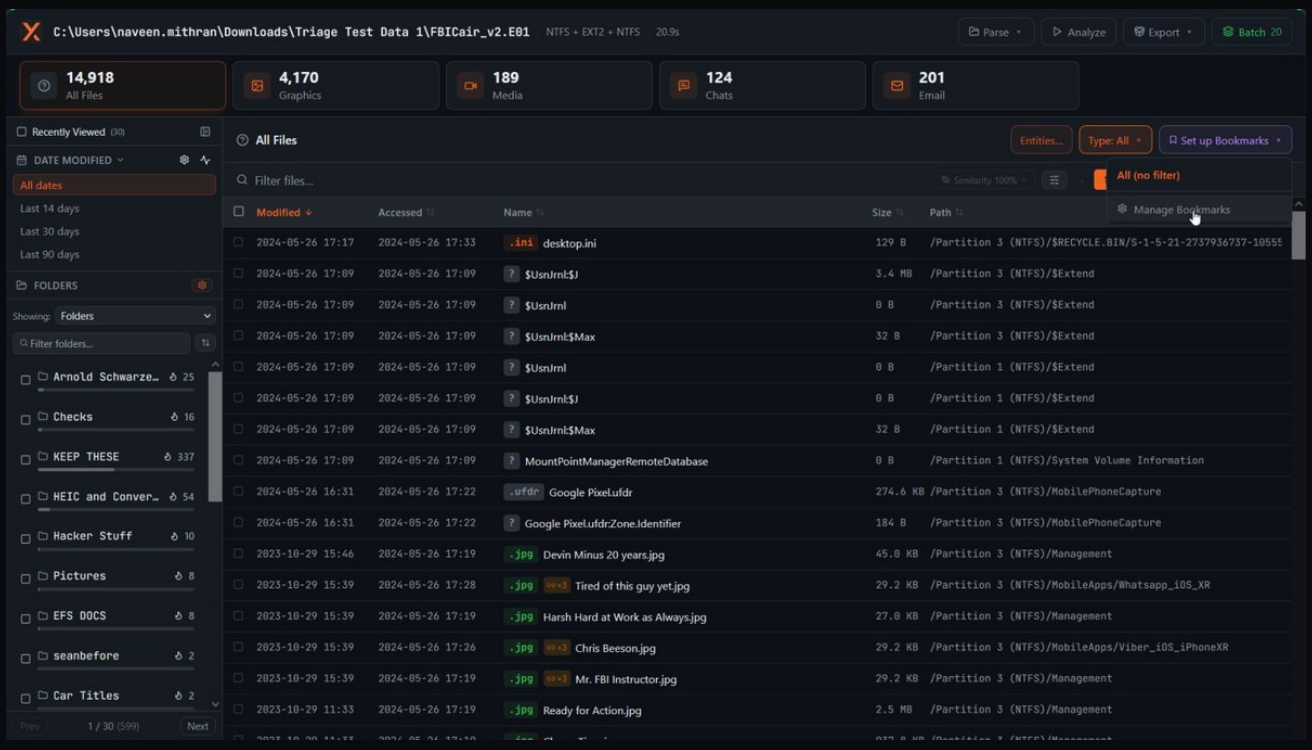

In the Manage Bookmarks dialog, give each bookmark a name, choose a color from 10 options, and assign a keyboard hotkey (Ctrl+0 through Ctrl+9). Up to 10 bookmarks can be created per case.

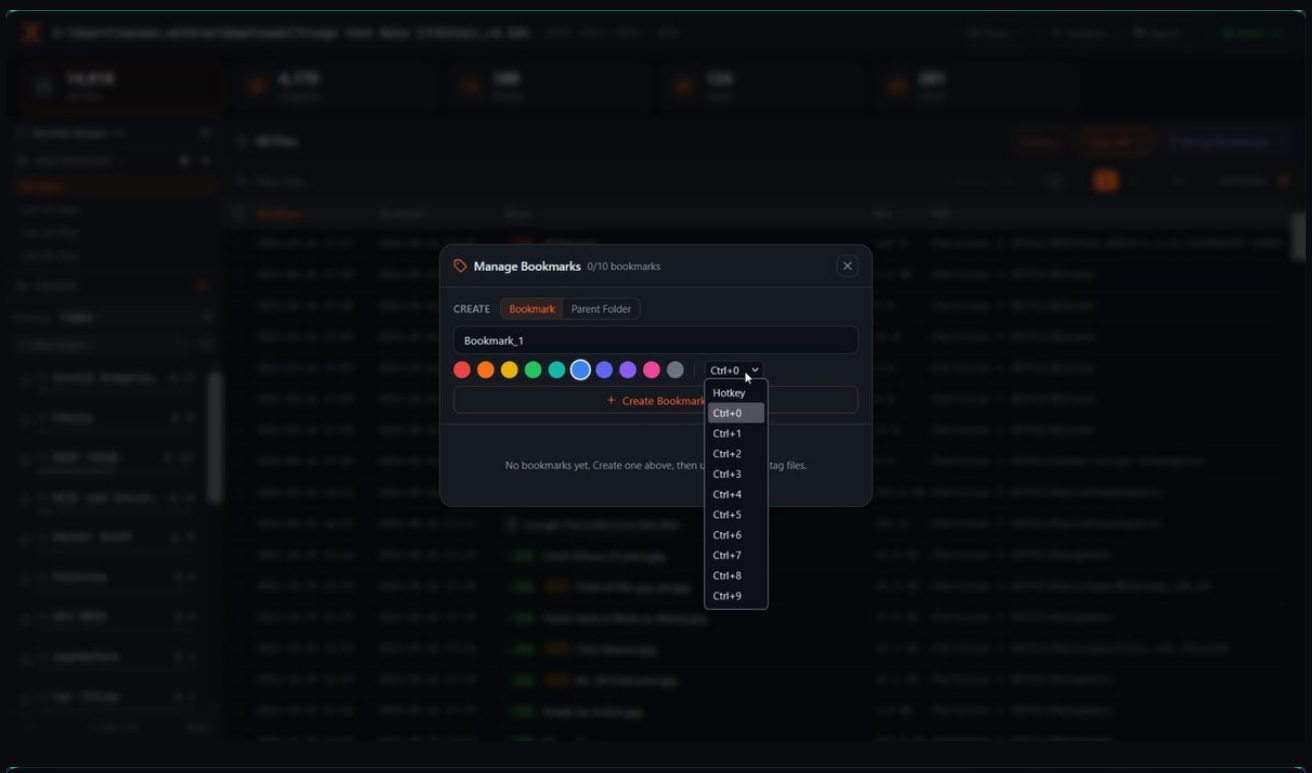

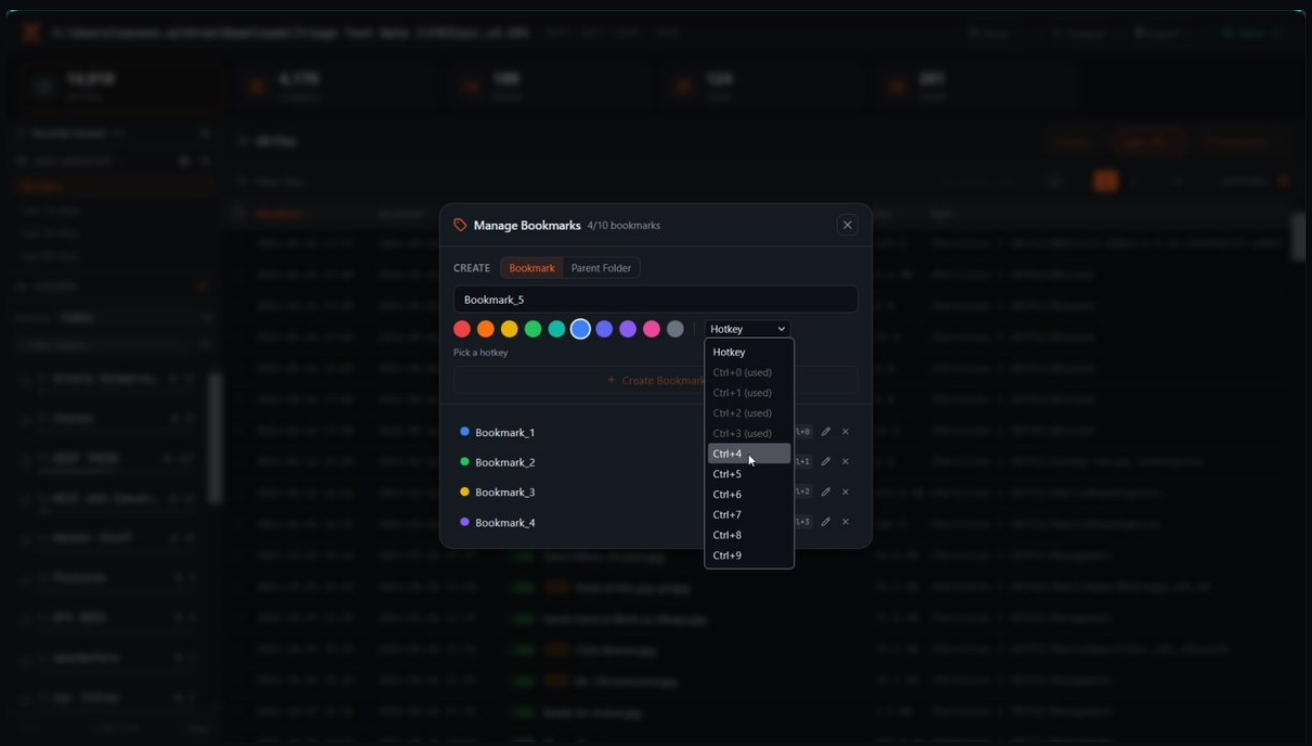

Plan your bookmark scheme before starting a review session. For example: blue for Key Evidence (Ctrl+0), green for Reviewed (Ctrl+1), yellow for Needs Follow-up (Ctrl+2).

Applying bookmarks

Select one or more files in the file list and press the hotkey for the bookmark you want to apply. The bookmark color dot appears immediately next to the filename.

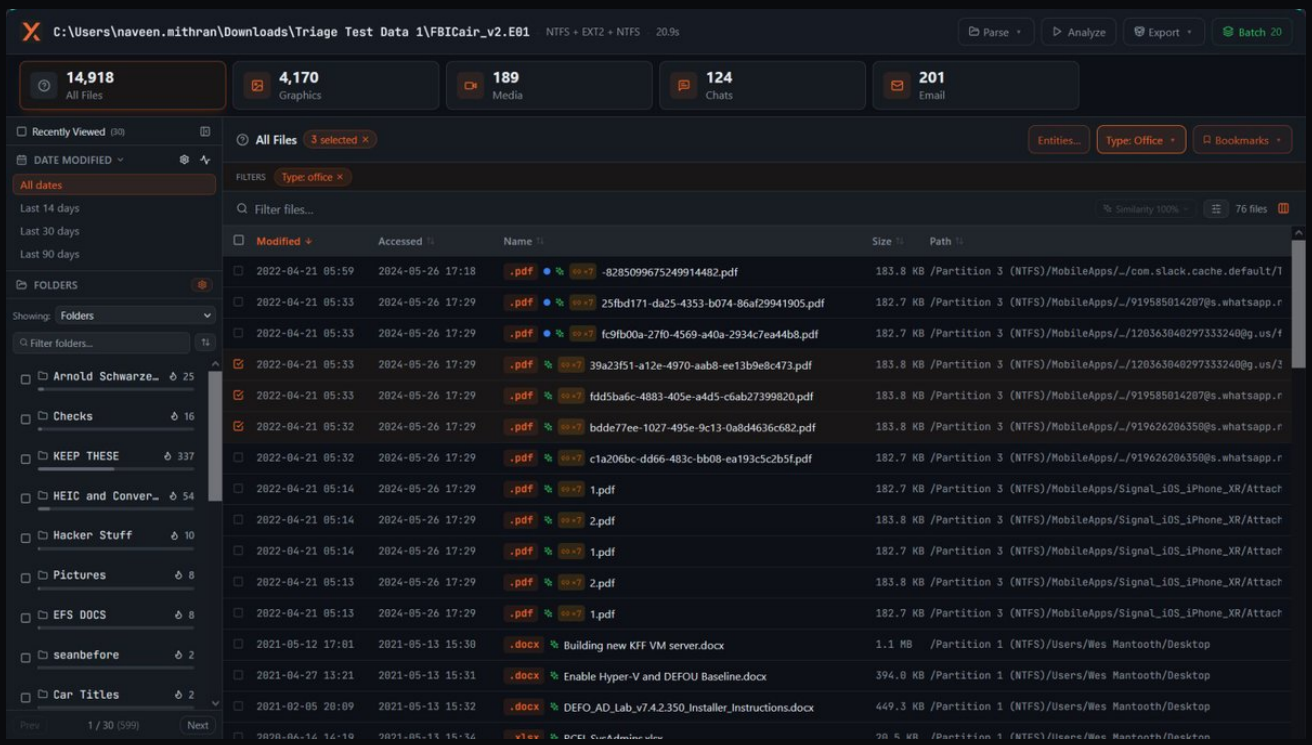

Multiple bookmarks can be applied to the same file. Each additional bookmark adds another colored dot next to the filename, letting you tag a file with multiple categories simultaneously.

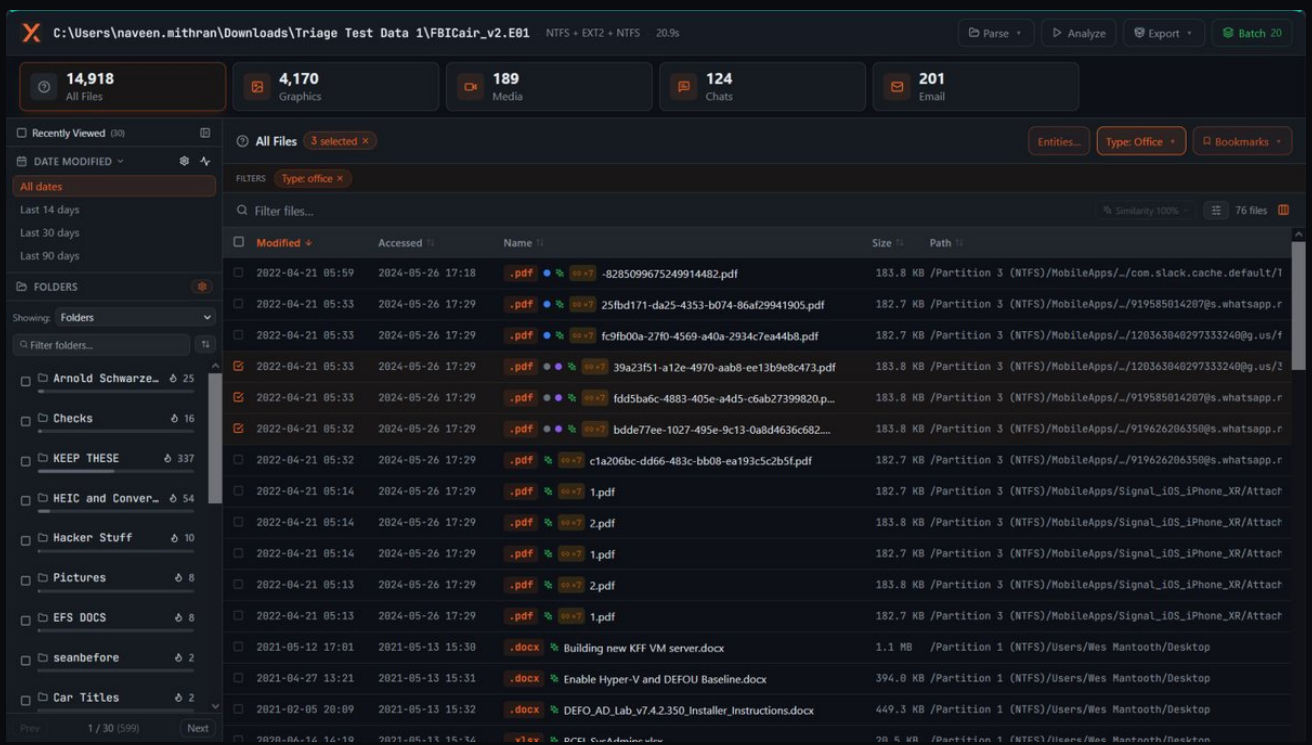

Filtering by Bookmarks

Once bookmarks are applied to files, use the Bookmarks filter to instantly surface all files tagged with a specific label across the entire file list.

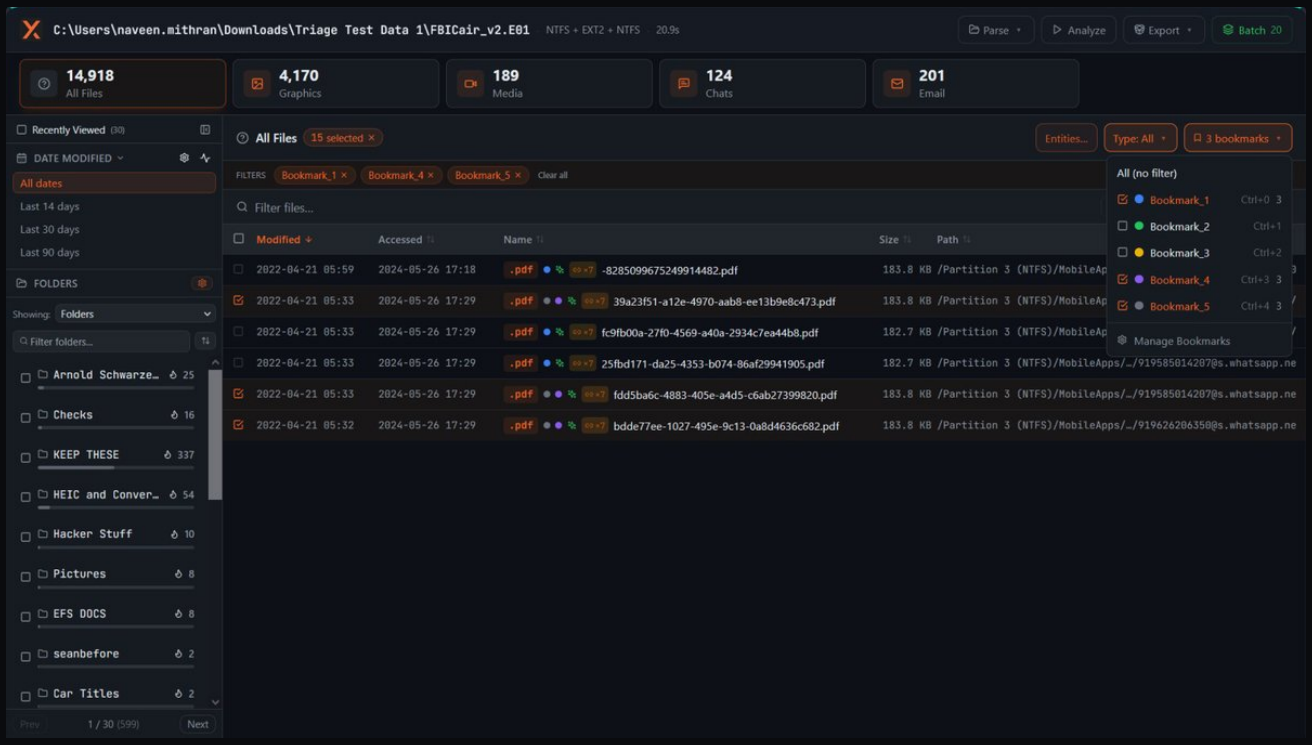

Click the Bookmarks button in the top-right toolbar to open the bookmark filter dropdown. Each bookmark is shown with its color dot, name, keyboard hotkey, and the number of files tagged with it. Click any bookmark to filter the file list instantly. To view all files regardless of bookmark, select All (no filter).

Multiple bookmark filters can be active at the same time. The active filters appear as colored badges in the Filters bar above the file list, and can be cleared individually or all at once.

Bookmarks persist for the entire duration of the case. Tagged files remain accessible across sessions, making bookmarks a reliable way to maintain a curated list of key evidence as the investigation progresses.

PART 12

EXPORT

Exporting native files, metadata, reports, and bookmarks to FTK / FTK Central

The Export section covers how to extract evidence and findings from ARMOURop into formats suitable for reporting, further analysis, or handoff to other forensic tools. Why this matters: Triage is only part of the workflow.

Export ensures that key evidence, search results, and AI-generated findings can be packaged and shared with prosecutors, moved into FTK, or archived as case documentation.

How To Access

Click the Export button in the top-right toolbar of the Triage Console to open the export dropdown.

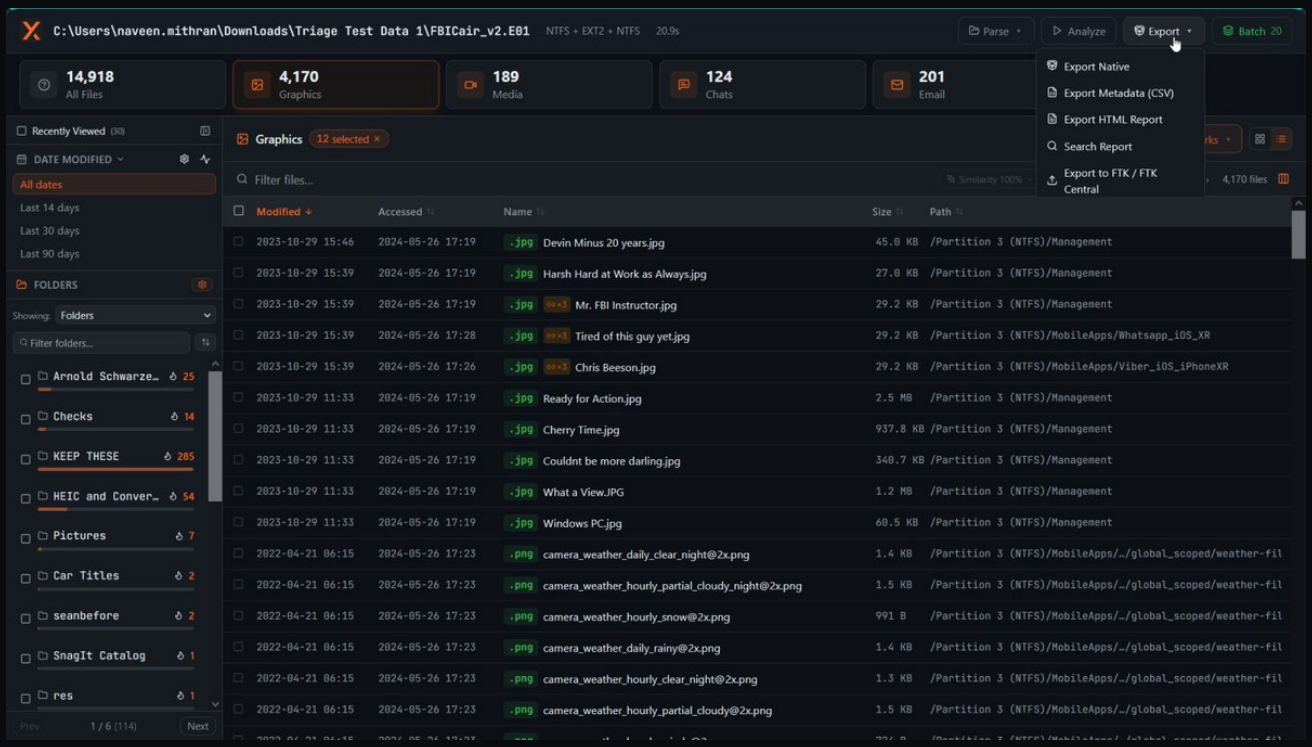

Export Native

Copies selected files to a destination folder in their original format, preserving the original directory structure.

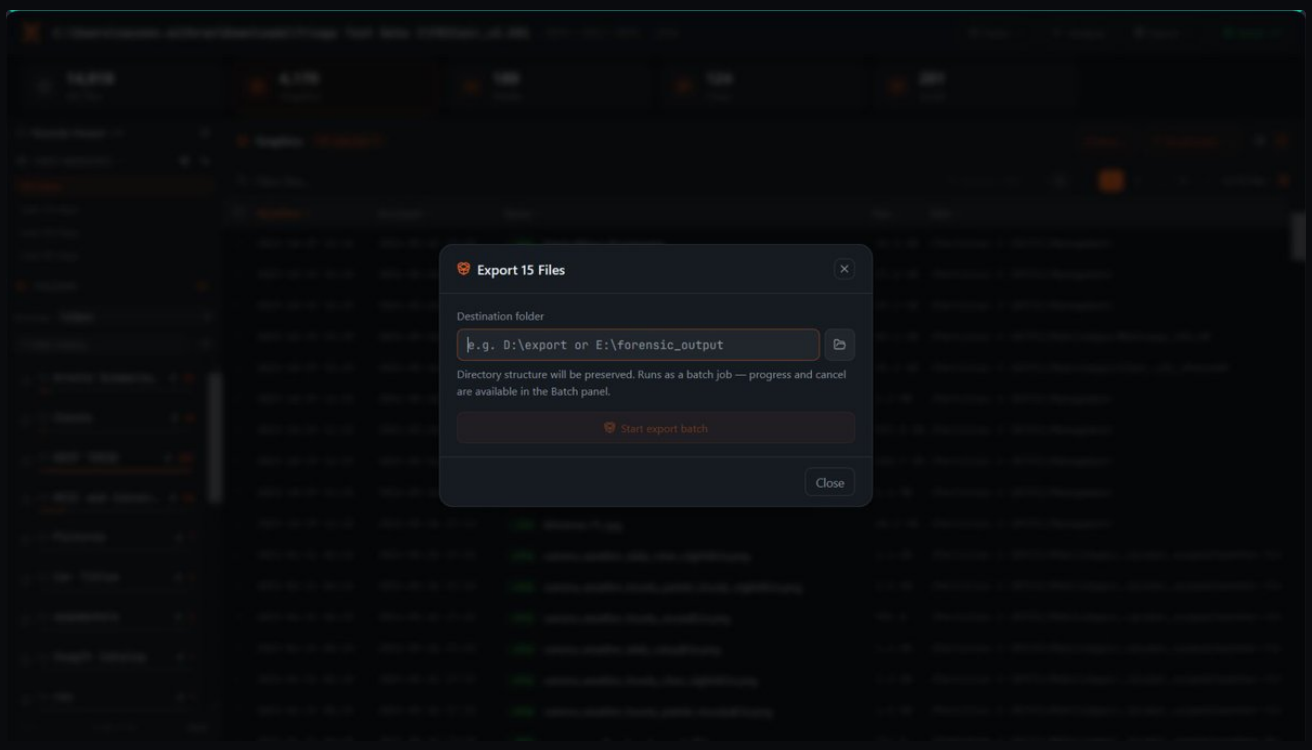

1. Select the files to export Use file list checkboxes, folder filters, or bookmark filters to select the files you want to export.

2. Enter the destination folder Type or browse to the output path (e.g. D:\export or E:\forensic_output).

3. Click Start export batch The export runs in the background. Monitor progress or cancel from the Batch panel in the top-right toolbar.

Export Metadata (CSV)

Exports metadata for the selected files as a CSV spreadsheet. The dialog shows a breakdown of the selection by file type and total size.

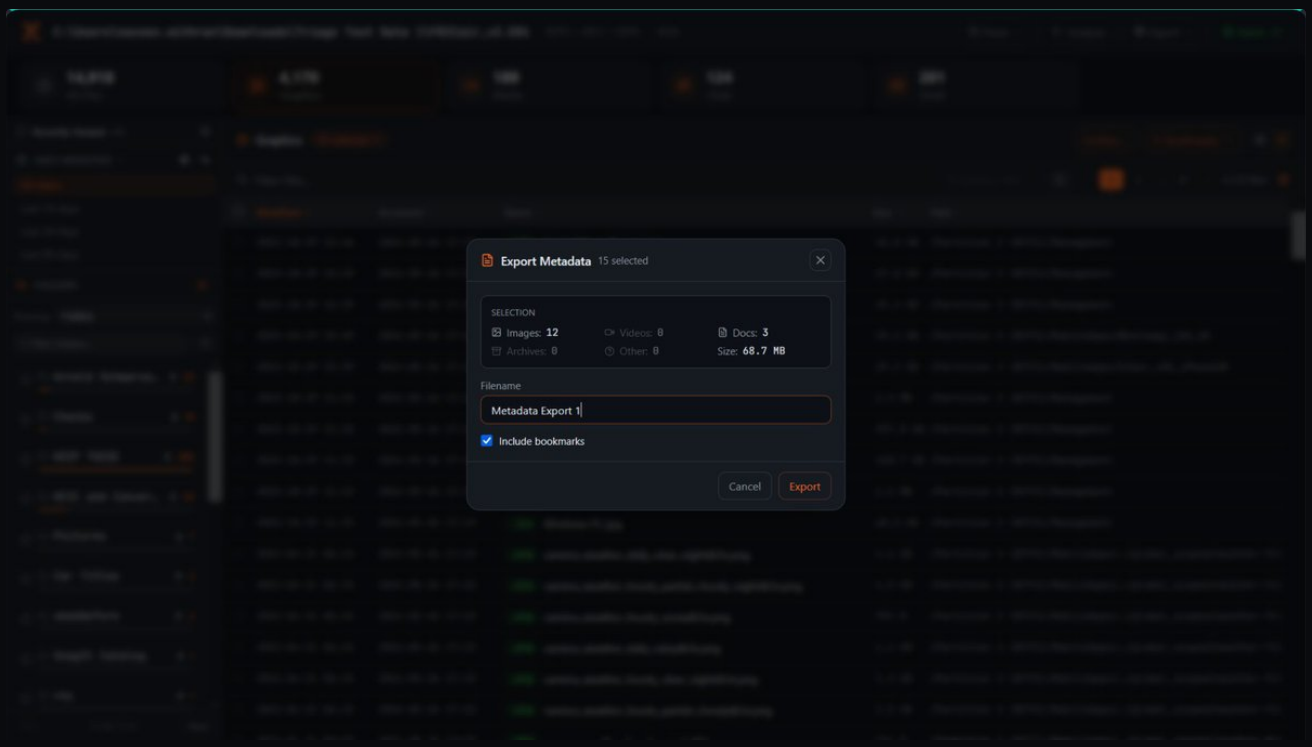

Give the export a filename, optionally include bookmark data, and click Export.

> The Include bookmarks option adds a column to the CSV showing which bookmarks are applied to each file. This makes it easy to sort and filter exported metadata by evidence category in Excel.

Export HTML Report

Generates a self-contained HTML report for the selected files. The report includes file details, thumbnails, AI-generated findings, and optionally bookmark labels.

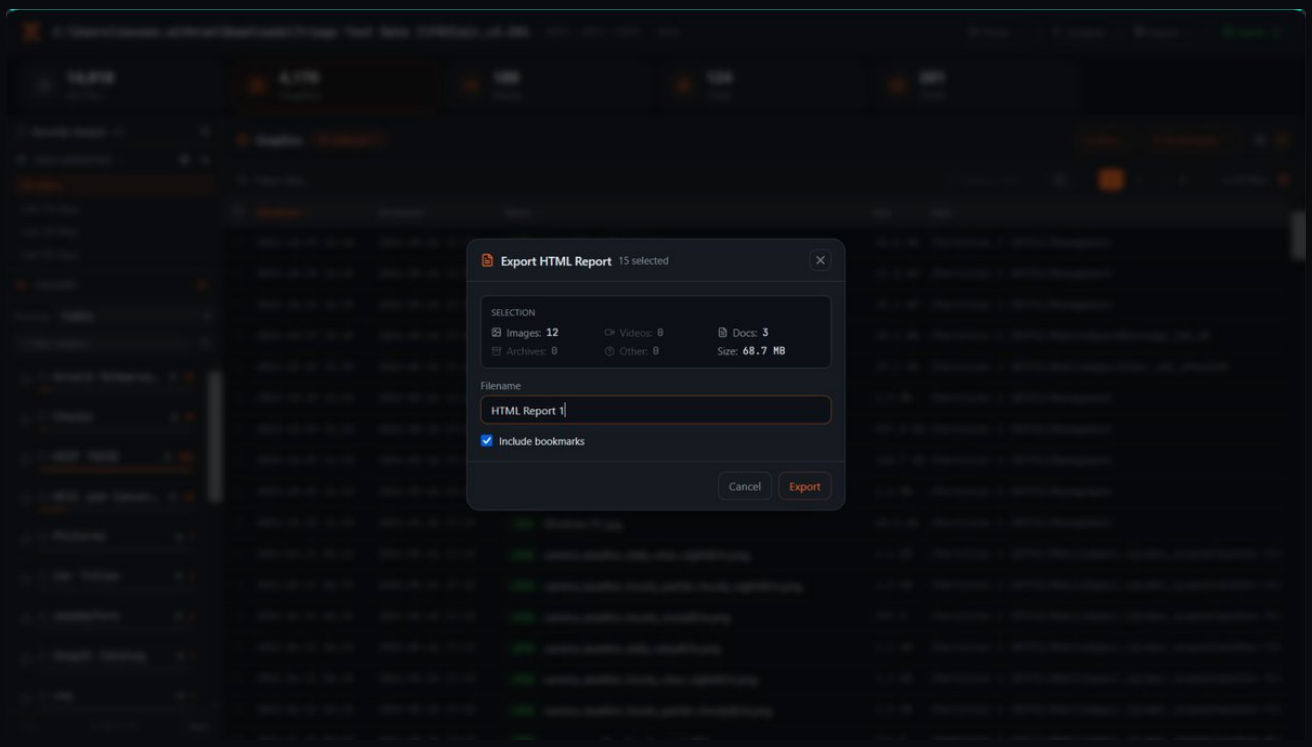

The report can be opened in any browser and shared without requiring ARMOURop to be installed.

Search Report

Runs a structured keyword search across all indexed documents and generates a report showing which files contain which keywords.

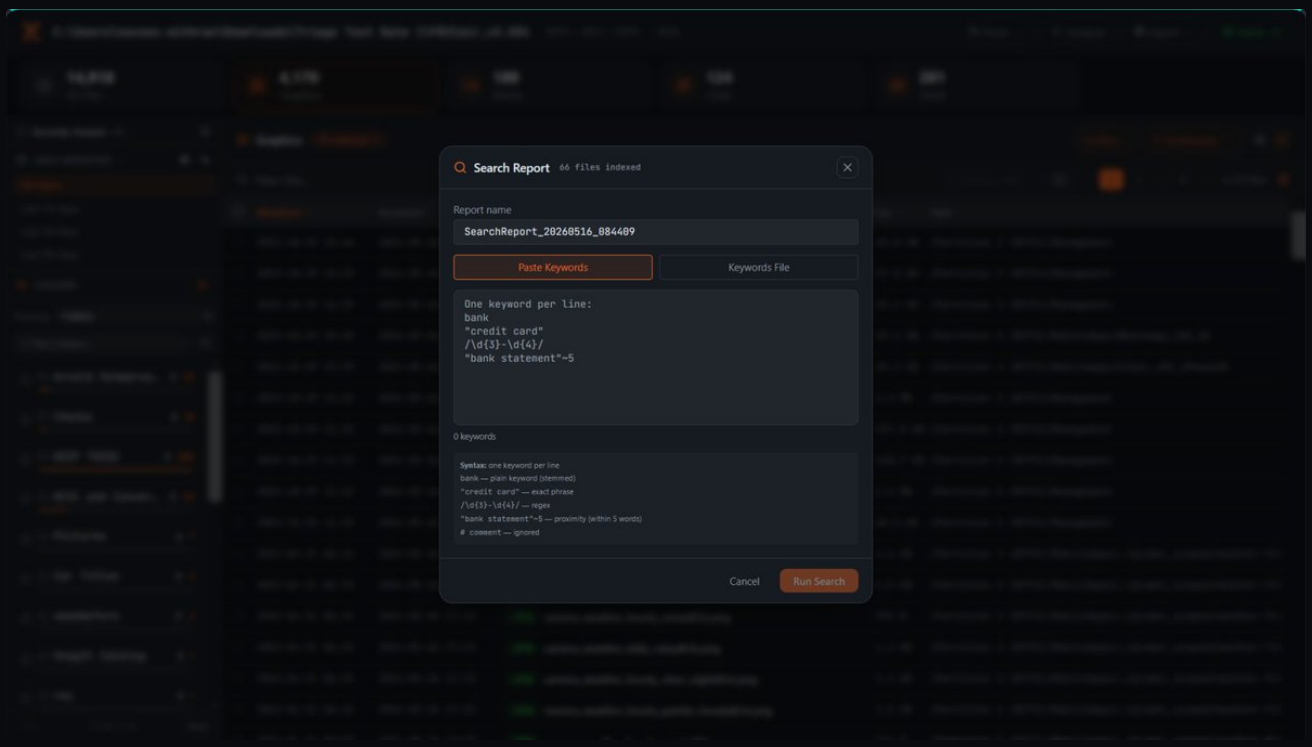

Keywords can be pasted directly or loaded from a file. The report is saved with a timestamped name that can be customized. Plain keyword e.g. bank - matches the word and stemmed variants. Exact phrase e.g. "credit card" - matches the exact phrase in quotes. Regex e.g. /\d{3}-\d{4}/ - matches patterns using regular expressions.

Proximity search e.g. "bank statement"~5 - words appearing within 5 words of each other. Comment Lines starting with # are ignored. Useful for annotating keyword lists.

Export to FTK / FTK Central

Exports bookmarked files directly into an existing FTK or FTK Central case. Each file is hashed (MD5) and matched to the corresponding object already processed in FTK/FTK Central. Prerequisite 1: Configure FTK connection The FTK/FTK Central Server URL and API Key must be configured in AI Settings before the export works.

Go to Settings and select the FTK/FTK Central tab to enter the server URL and API key. Prerequisite 2: Evidence in FTK The same evidence must already be processed in an FTK or FTK Central case. ARMOURop matches bookmarked files by hash to objects in that case.

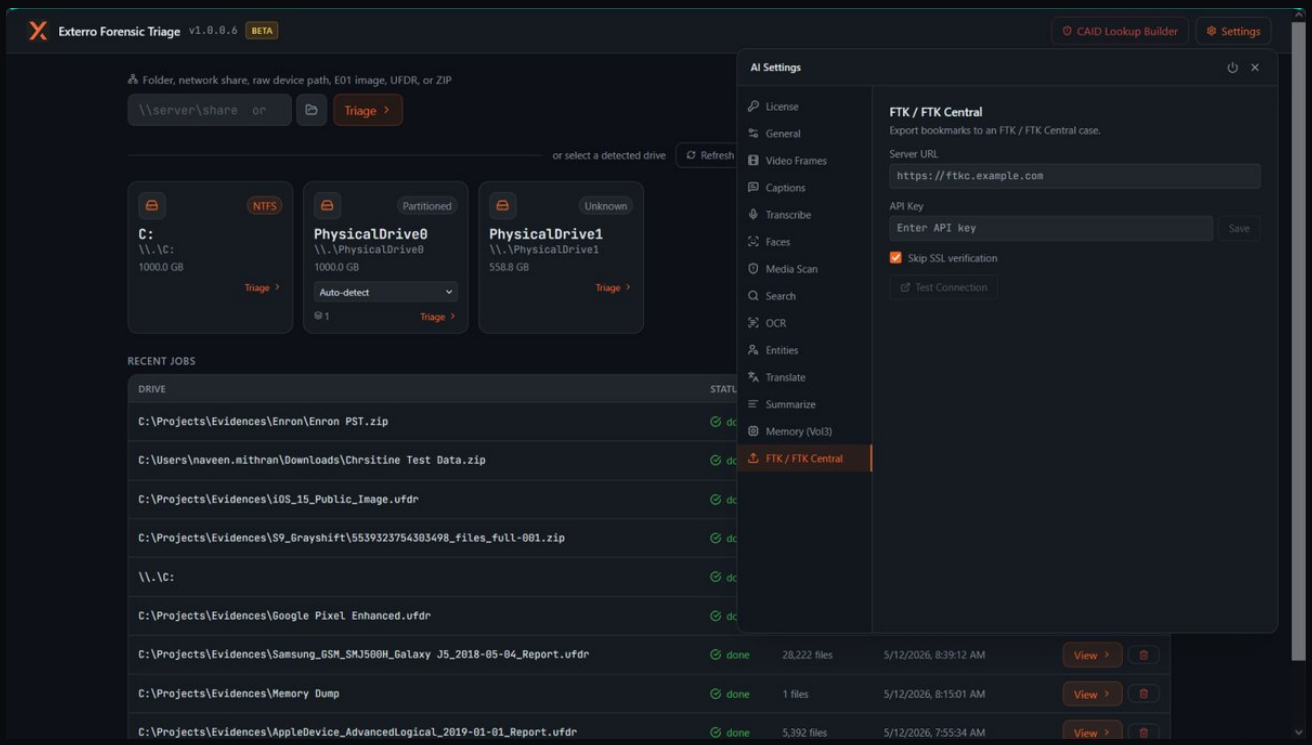

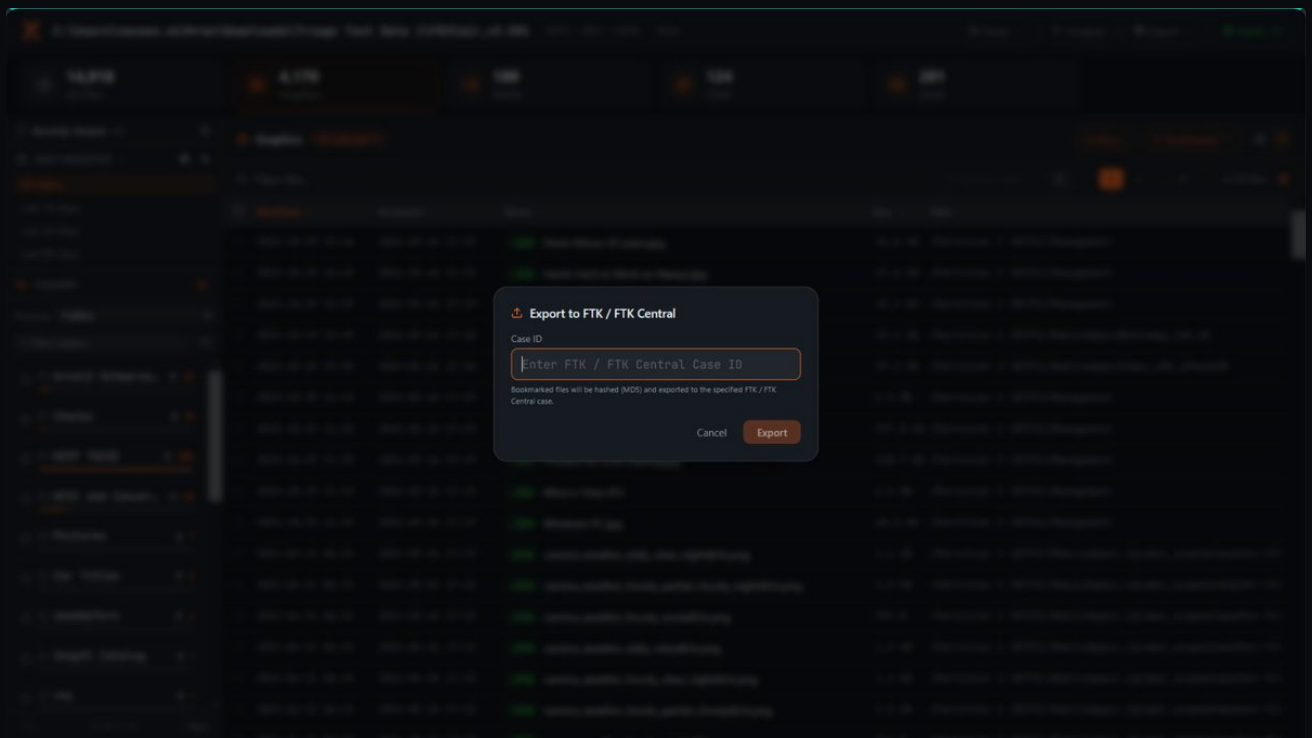

1. Configure the FTK connection in Settings Go to Settings, select FTK/FTK Central, enter the Server URL and API Key, then click Test Connection.

2. Ensure evidence is processed in FTK The same evidence must exist in an FTK or FTK Central case. Note the Case ID you want to export into.

3. Apply bookmarks to files in ARMOURop Only bookmarked files are exported. Make sure all relevant files are tagged with at least one bookmark.

4. Enter the Case ID and export Select Export to FTK/FTK Central, enter the Case ID, and click Export.

PART 13

SETTINGS

Configuring the AI engine: licensing, models, thresholds, and integrations

The AI Settings panel is where you configure the AI engine that powers ARMOURop. From here you activate your license, choose which AI models run for each feature, tune performance for your hardware, and set up integrations with external tools.

How To Access

Click the Settings button in the top-right corner of the Home screen to open the AI Settings panel. The panel is organised into 14 sections, each covering a specific part of the AI engine. Use the left sidebar within the panel to navigate between sections. License Activate your license key or check license status.

General Configure the Ollama/cloud endpoint, CPU thread allocation, small-file filters, GPU detection, AI model storage path, and global bookmark templates. Video Frames Control how video frames are extracted for AI analysis. Captions Choose the model used to generate image and video captions. Transcribe Choose the Whisper model used for speech-to-text transcription.

Faces Configure the face detection model and minimum face size. Media Scan Configure the Quick Media Scan model and manage detection categories. Search Choose the semantic search embedding model and set confidence thresholds. OCR Configure the OCR text extraction model. Entities Choose the model used to extract named entities from documents and chats.

Translate Choose the translation model for Translate to English. Summarize Choose the model used to generate document and email summaries. Memory (Vol3) Configure YARA rules for memory analysis via Volatility 3. FTK / FTK Central Set up the connection for exporting bookmarks to FTK or FTK Central cases.

License

Activate ARMOURop by entering your license key. The License section also shows your Machine ID, which is needed when requesting a license from Exterro.

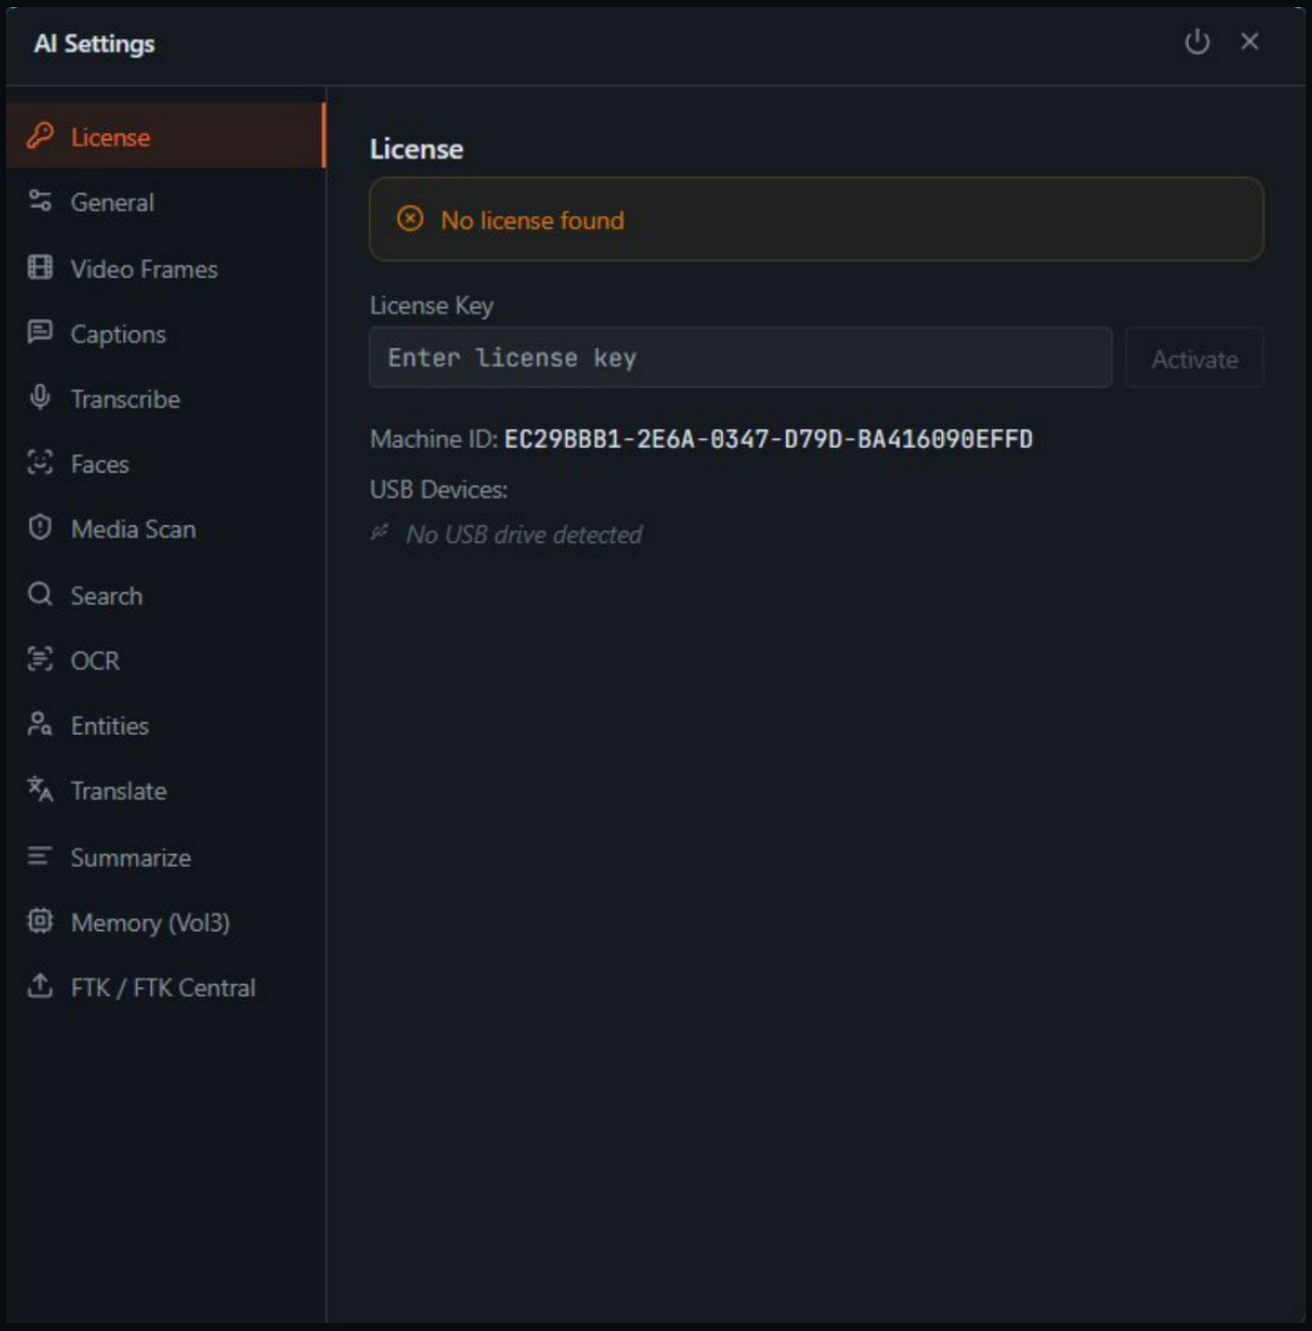

1. Enter your license key Paste the license key provided by Exterro into the License Key field.

2. Click Activate Click the Activate button to validate and activate the license.

3. Verify activation Once activated, the "No license found" warning banner will be replaced with your license details. MACHINE ID AND USB LICENSING Your Machine ID is displayed below the license key field. This is a unique identifier tied to the machine ARMOURop is installed on. Provide this ID to Exterro when requesting or renewing a license. The USB Devices section shows whether a hardware license key (USB dongle) is connected. If USB licensing is used, the dongle must be inserted before launching ARMOURop.

> If you are running ARMOURop from a USB drive on a live machine, the Machine ID will reflect the host machine, not the USB drive.

General

General settings control the core AI engine configuration: compute resource allocation, Ollama and cloud model endpoints, file filtering defaults, GPU detection, and the global bookmark template.

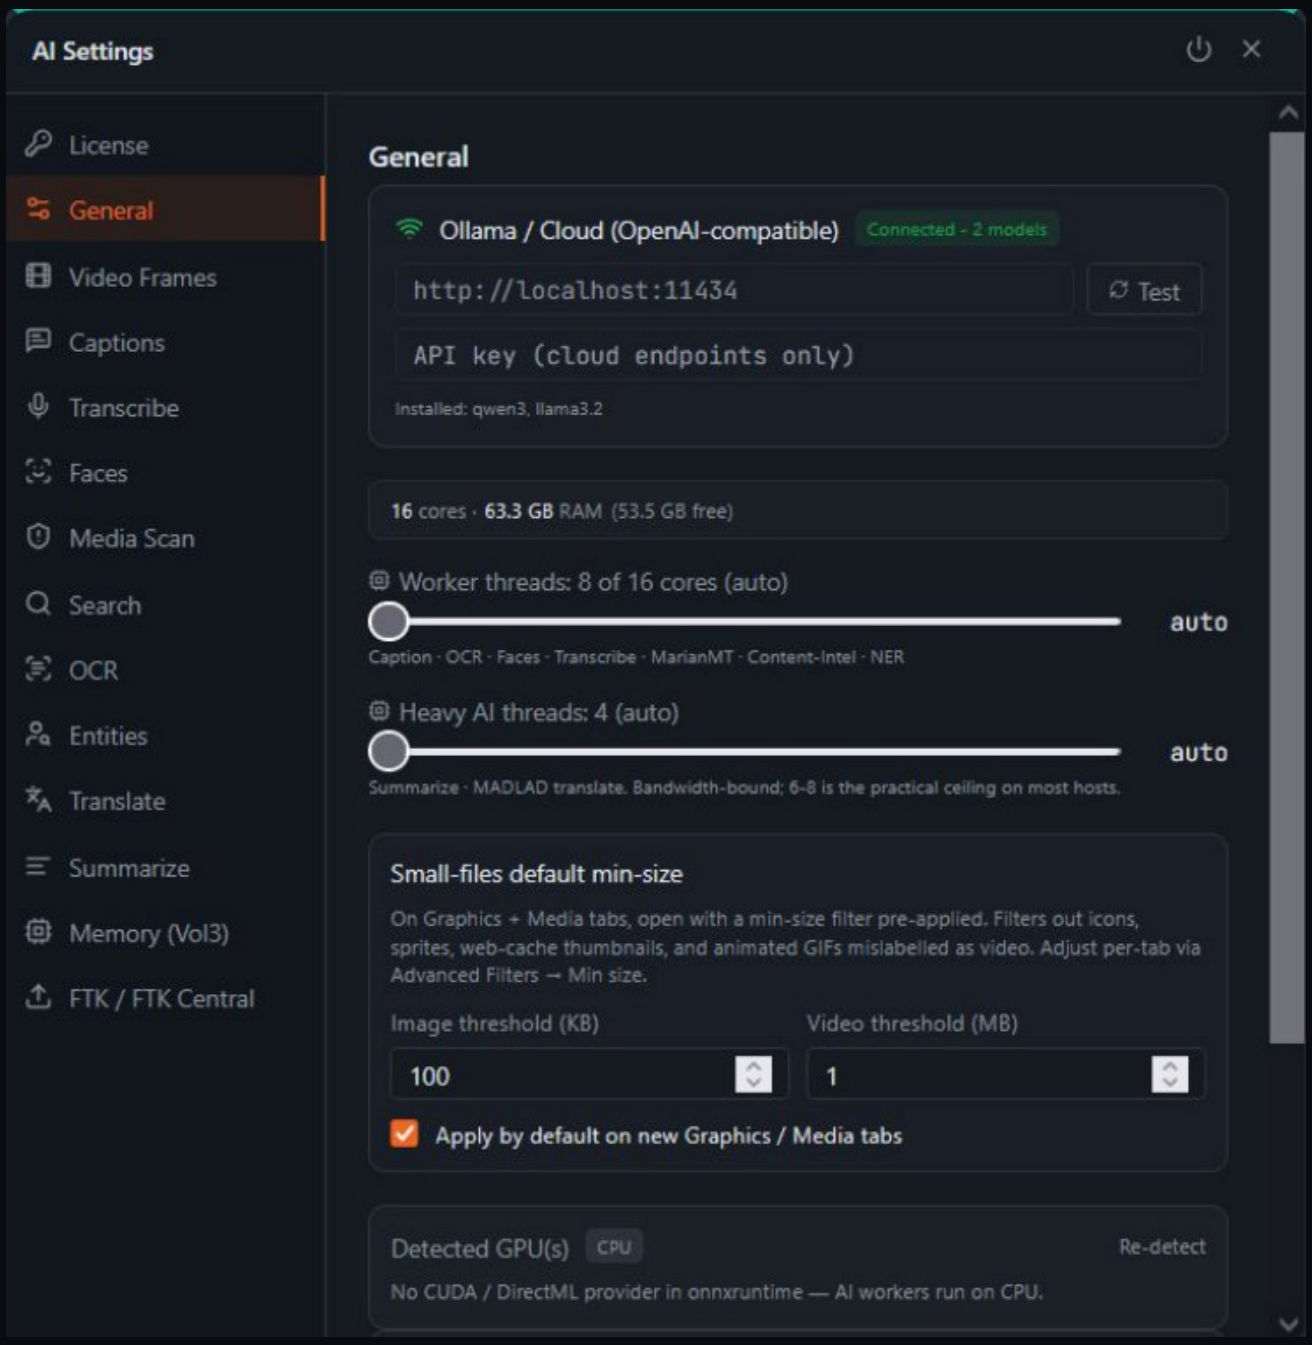

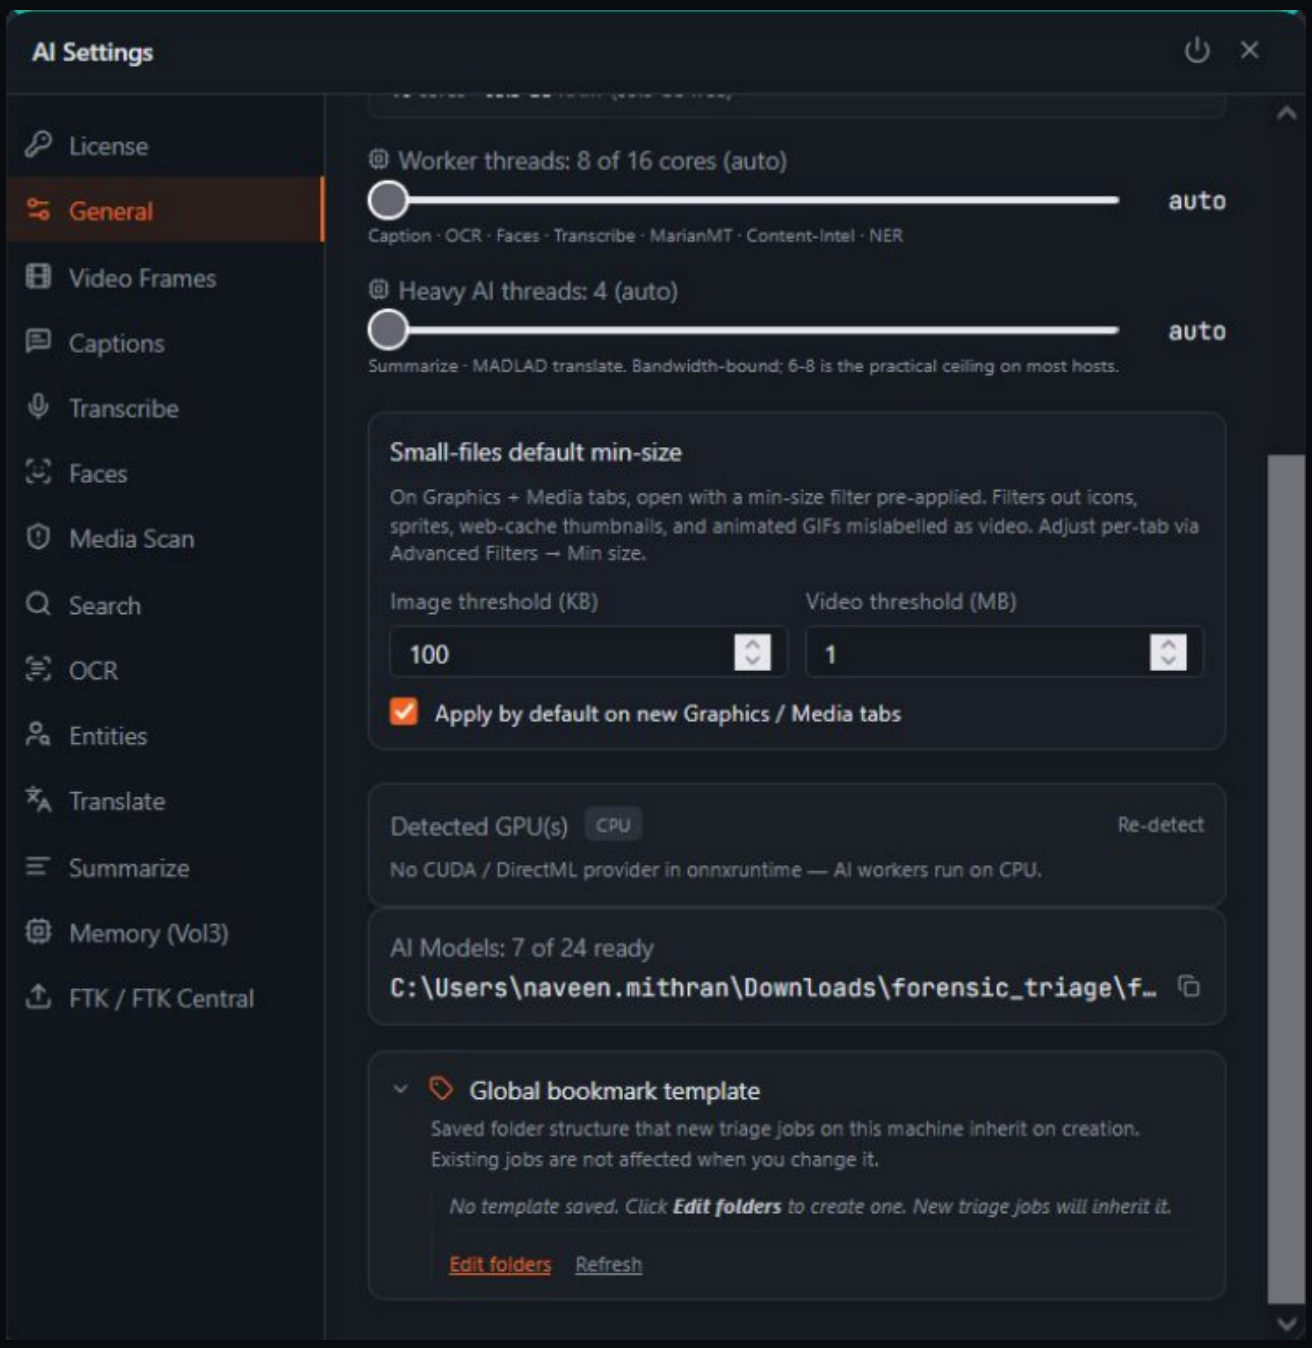

Ollama / Cloud (Openai-Compatible)

This section connects ARMOURop to a local Ollama instance or any OpenAI-compatible cloud endpoint. The connection is used by AI features that support Ollama or cloud models, such as Captions, Summarize, and Semantic Search. Enter the server URL and optionally an API key for cloud endpoints, then click Test to verify the connection.

> When connected, the available models installed on the Ollama instance are shown (e.g. qwen3, llama3.2). These models then appear as selectable options in the relevant feature settings.

Worker Threads

Controls how many CPU cores are allocated to lighter AI tasks: Caption, OCR, Faces, Transcribe, MarianMT, Content-Intel, and NER. The default is auto, which lets the system allocate cores based on available hardware. Drag the slider to set a manual value if you want to reserve cores for other work on the machine.

Heavy Ai Threads

Controls CPU cores allocated to heavier AI tasks: Summarize and MADLAD translation. The default is auto. The UI notes that 6 to 8 threads is the practical ceiling on most hosts, beyond which returns diminish.

Small-Files Default Min-Size

Pre-applies a minimum file size filter on the Graphics and Media tabs when a new job is opened. This filters out icons, sprites, web-cache thumbnails, and animated GIFs that are mislabelled as video. The default thresholds are 100KB for images and 1MB for video. These can be adjusted per-tab via Advanced Filters after opening a job.

Detected Gpus

Shows whether a CUDA or DirectML GPU is available for AI acceleration. If no compatible GPU is found, the tool runs all AI tasks on CPU. Click Re-detect after installing GPU drivers or connecting a GPU to refresh this detection.

Ai Models

Shows how many of the available AI models are downloaded and ready (e.g. 7 of 24 ready). The model storage path is shown below. Models not yet downloaded show a Get button in their respective settings sections.

Global Bookmark Template

Defines a bookmark folder structure that all new triage jobs on this machine inherit automatically when created. This is useful for standardising case organisation across investigations. Existing jobs are not affected when the template is changed. Click Edit folders to create or modify the template.

AI Assistant Settings

Configures the case chat assistant (the Assistant button on a case) and the local Ollama model it talks to.

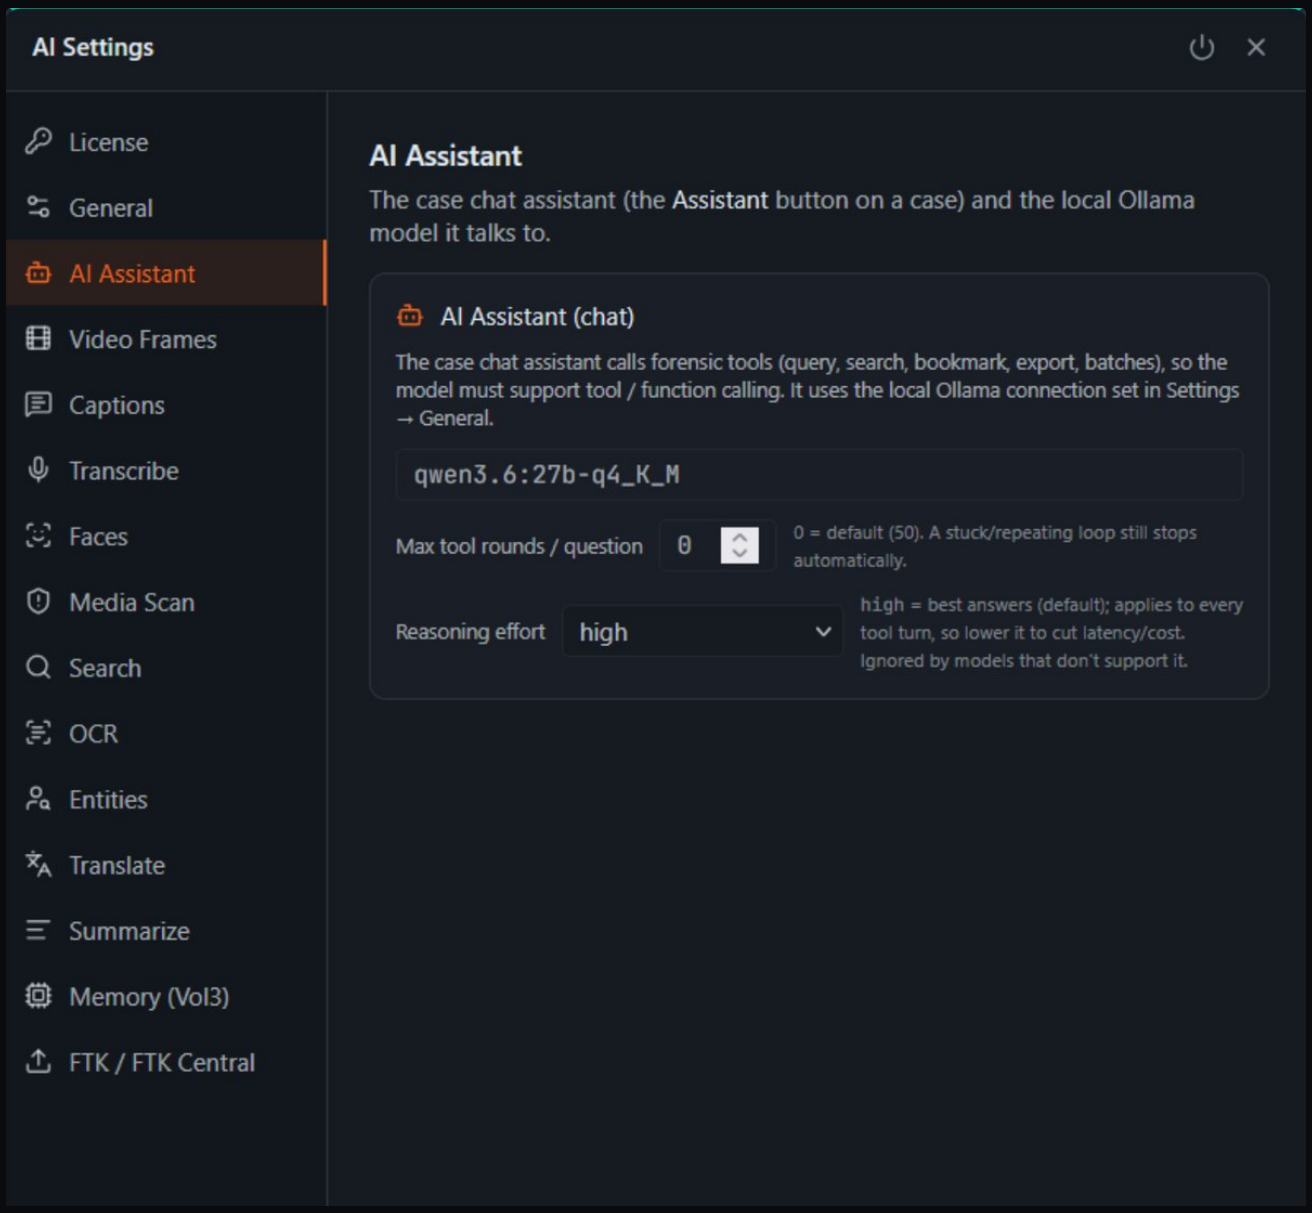

Model The Ollama model tag the assistant uses (e.g. qwen3.6:27b-q4_K_M). The case chat assistant calls forensic tools (query, search, bookmark, export, batches), so the model must support tool / function calling. It uses the local Ollama connection configured in Settings → General.

Max tool rounds / question Caps how many tool calls the assistant can chain together while answering a single question. 0 uses the default of 50. A stuck or repeating tool-call loop still stops automatically regardless of this setting.

Reasoning effort high (default) gives the best answers but applies to every tool-calling turn, so lower it to cut latency and cost on longer conversations. Ignored by models that don't support a reasoning-effort parameter.

> The AI Assistant requires a tool-calling capable model. If responses fail to use case tools (query, search, bookmark, export, batches), confirm the selected model supports function calling and that the Ollama connection in Settings → General is reachable.

Video Frames

Controls how frames are extracted from videos for AI analysis, and sets the minimum image dimension filter applied across all AI features.

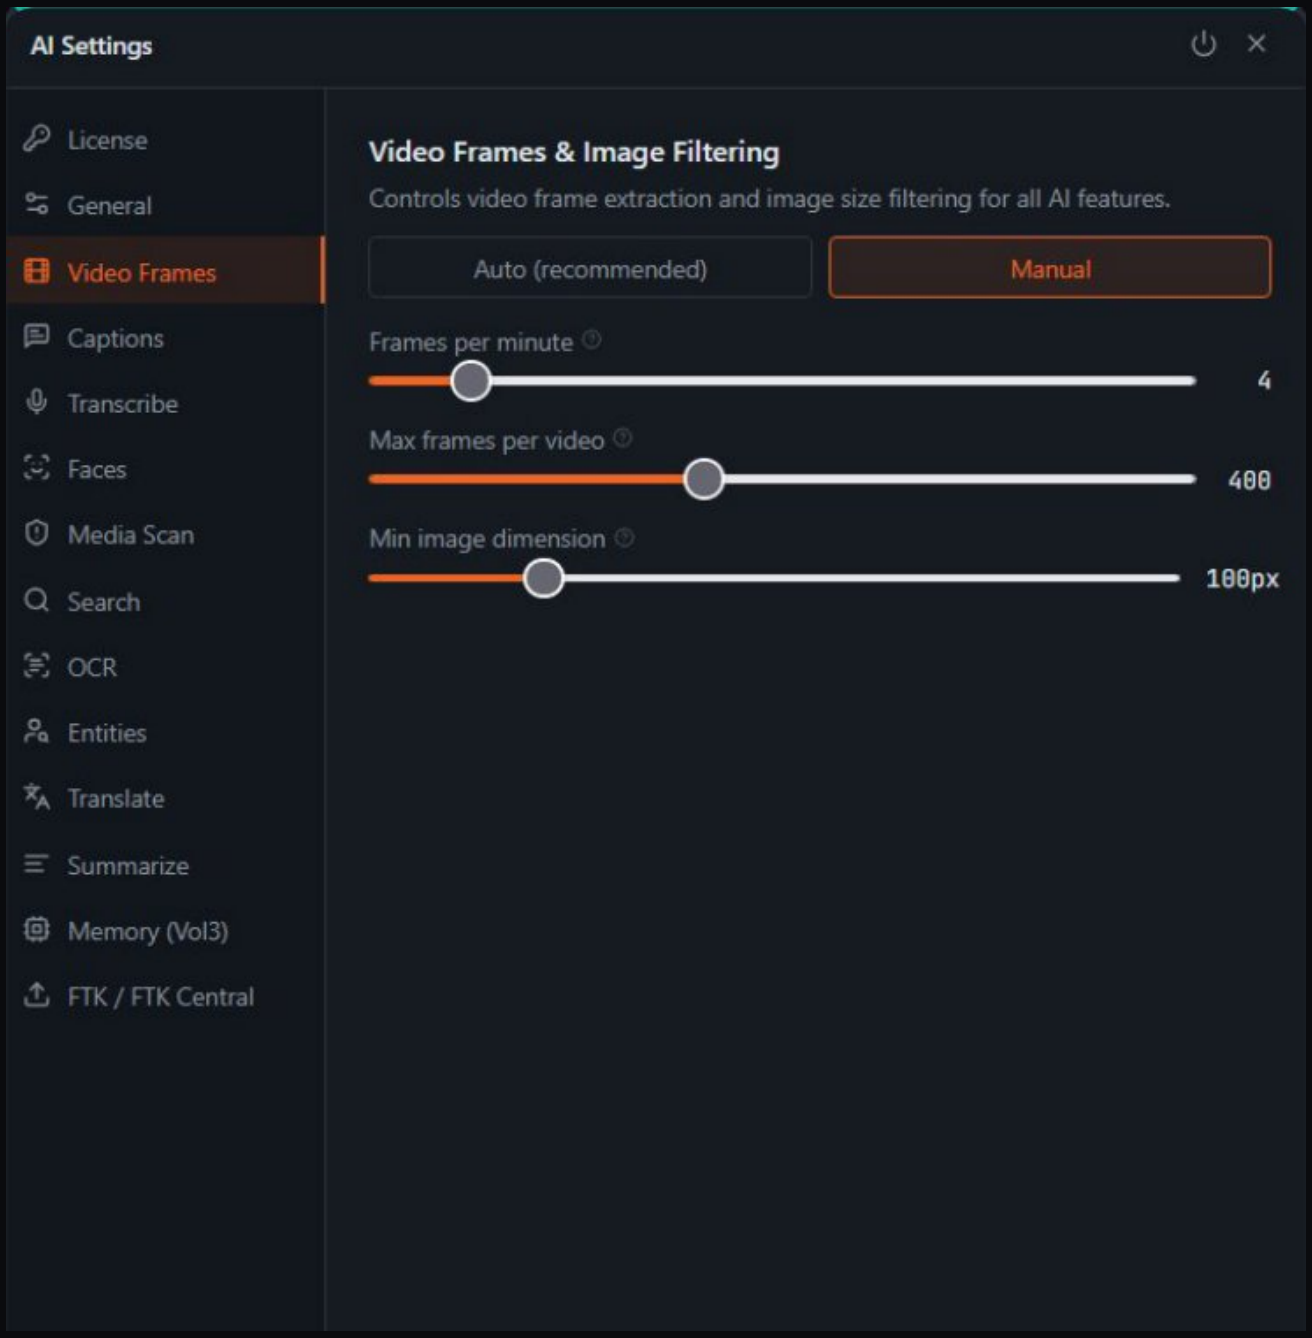

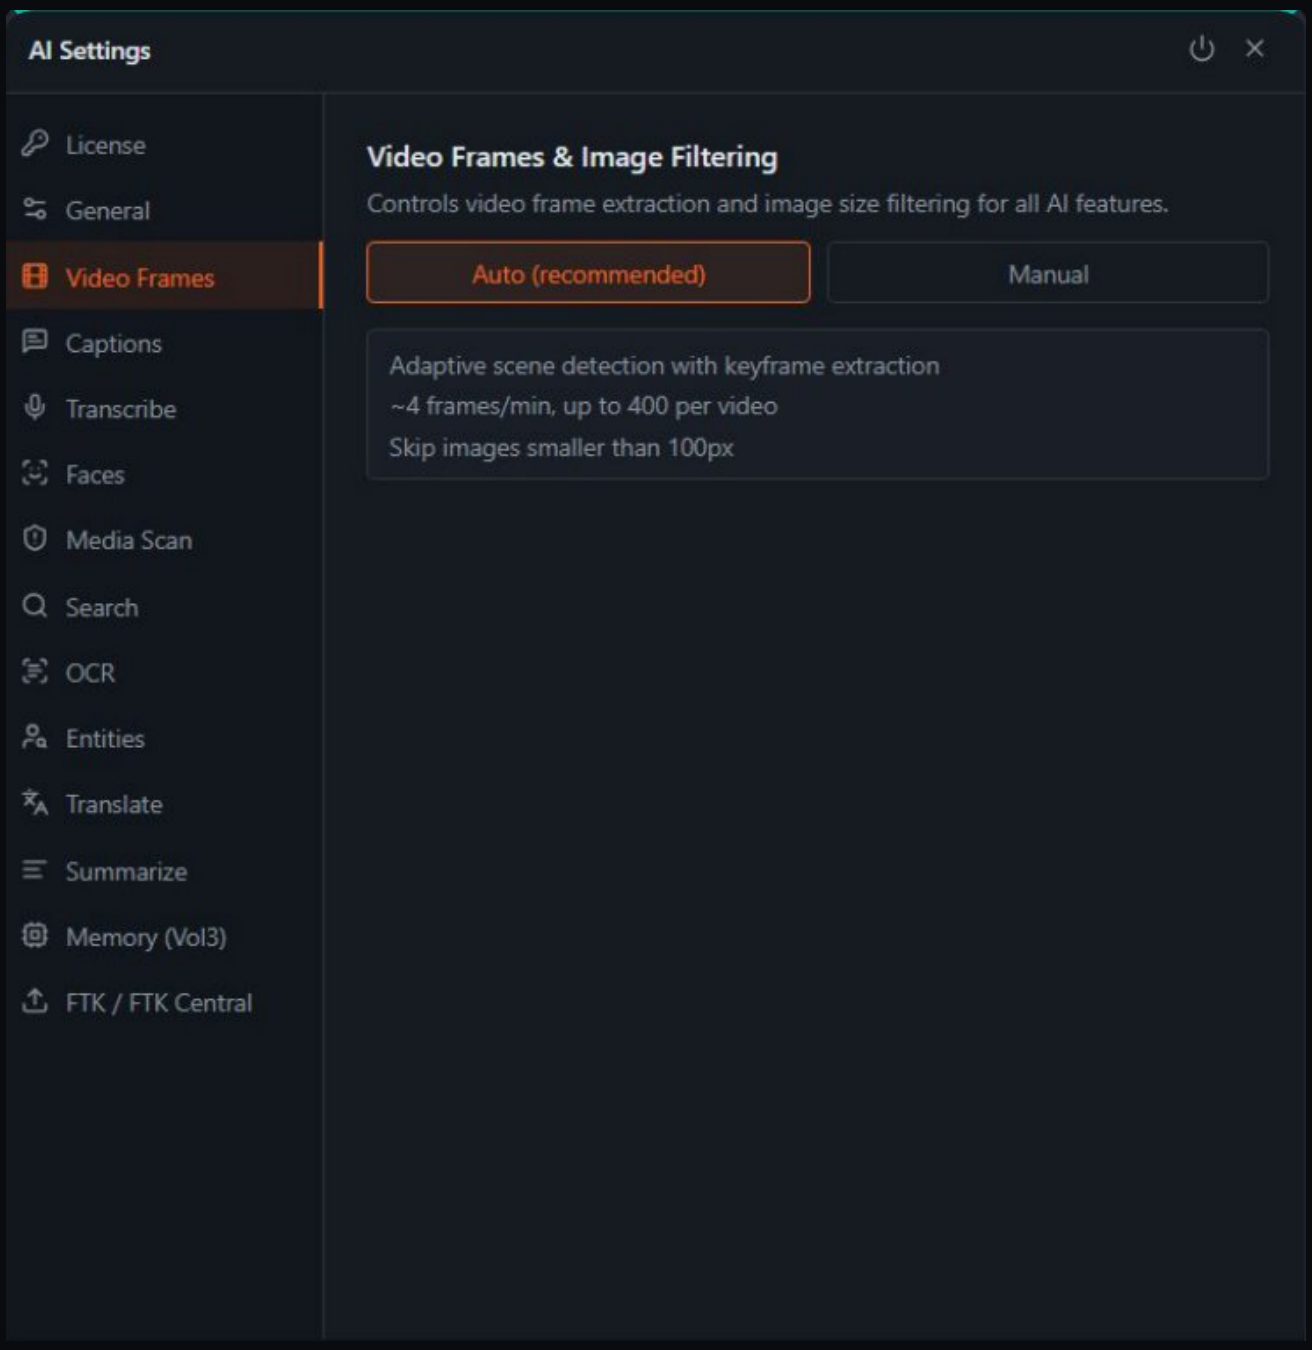

Auto mode (recommended) In Auto mode, the tool uses adaptive scene detection with keyframe extraction. This intelligently identifies meaningful scene changes rather than extracting frames at fixed intervals, resulting in more representative frame coverage.

The defaults are approximately 4 frames per minute, a maximum of 400 frames per video, and images smaller than 100px are skipped.

Auto mode is recommended for most investigations. It balances coverage and performance without requiring manual tuning. Manual mode Manual mode exposes three sliders for precise control: Frames per minute How many frames to extract per minute of video. Default: 4. Higher values give more coverage but increase processing time and storage.

Max frames per video Hard cap on the total number of frames extracted per video regardless of duration. Default: 400. Min image dimension Skips any image or extracted frame smaller than this pixel dimension. Default: 100px. Increase this to filter out small thumbnails and icons.

Caption Settings

Configures the AI model used to generate text descriptions for images and video frames. Captions power in-image and in-video keyword search across the Graphics and Media tabs.

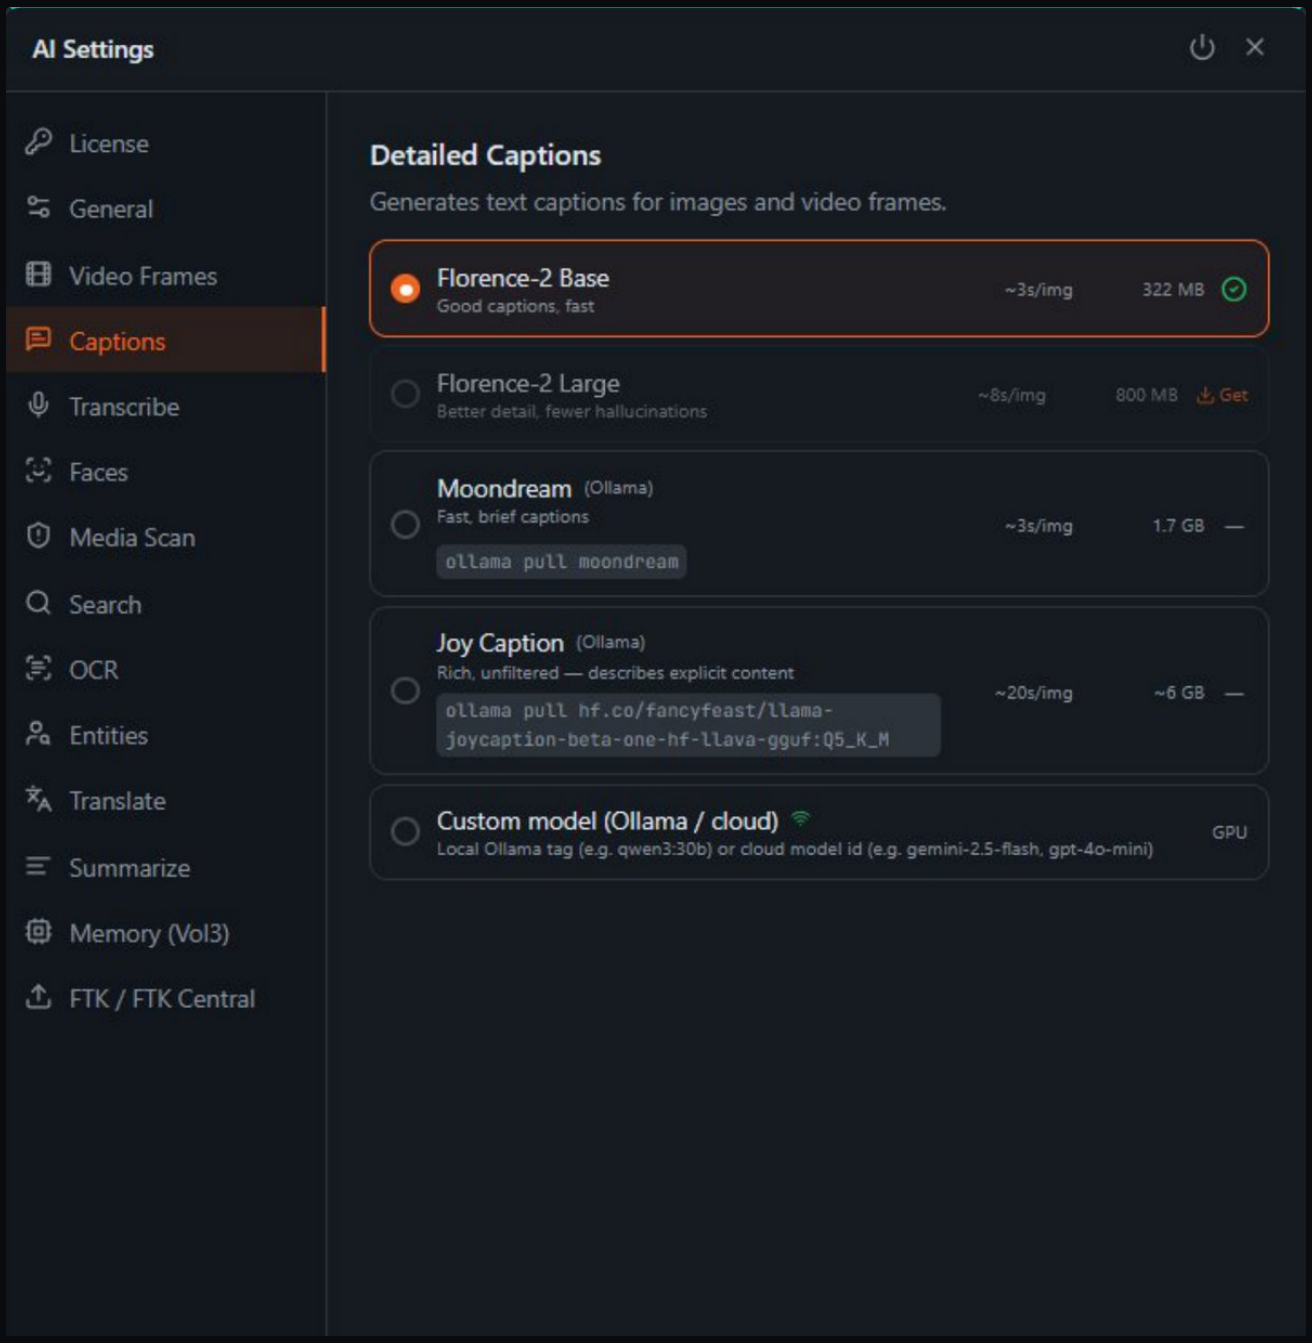

Florence-2 Base

Default

Good captions, fast. Best starting point for most investigations. Florence-2 Large Better detail with fewer hallucinations. Worth downloading for cases where image content needs more precise descriptions. Moondream

Ollama

Fast, brief captions. Requires Ollama. Install via ollama pull moondream . Joy Caption

Ollama

Rich, unfiltered captions. Describes explicit content in detail. Requires Ollama. Significantly slower and larger than other options. Custom model (Ollama / cloud) Use any local Ollama tag or cloud model ID (e.g. gemini-2.5-flash, gpt-4o-mini). Requires the Ollama/cloud endpoint to be configured in General settings.

> Florence-2 Base is downloaded and ready by default. For most investigations this is sufficient. Florence-2 Large offers better results on complex or ambiguous images if you have time to download it.

Transcribe

Configures the Whisper model used to convert speech in audio and video files to searchable text. Always runs locally. The right model depends on the audio quality and language of your evidence.

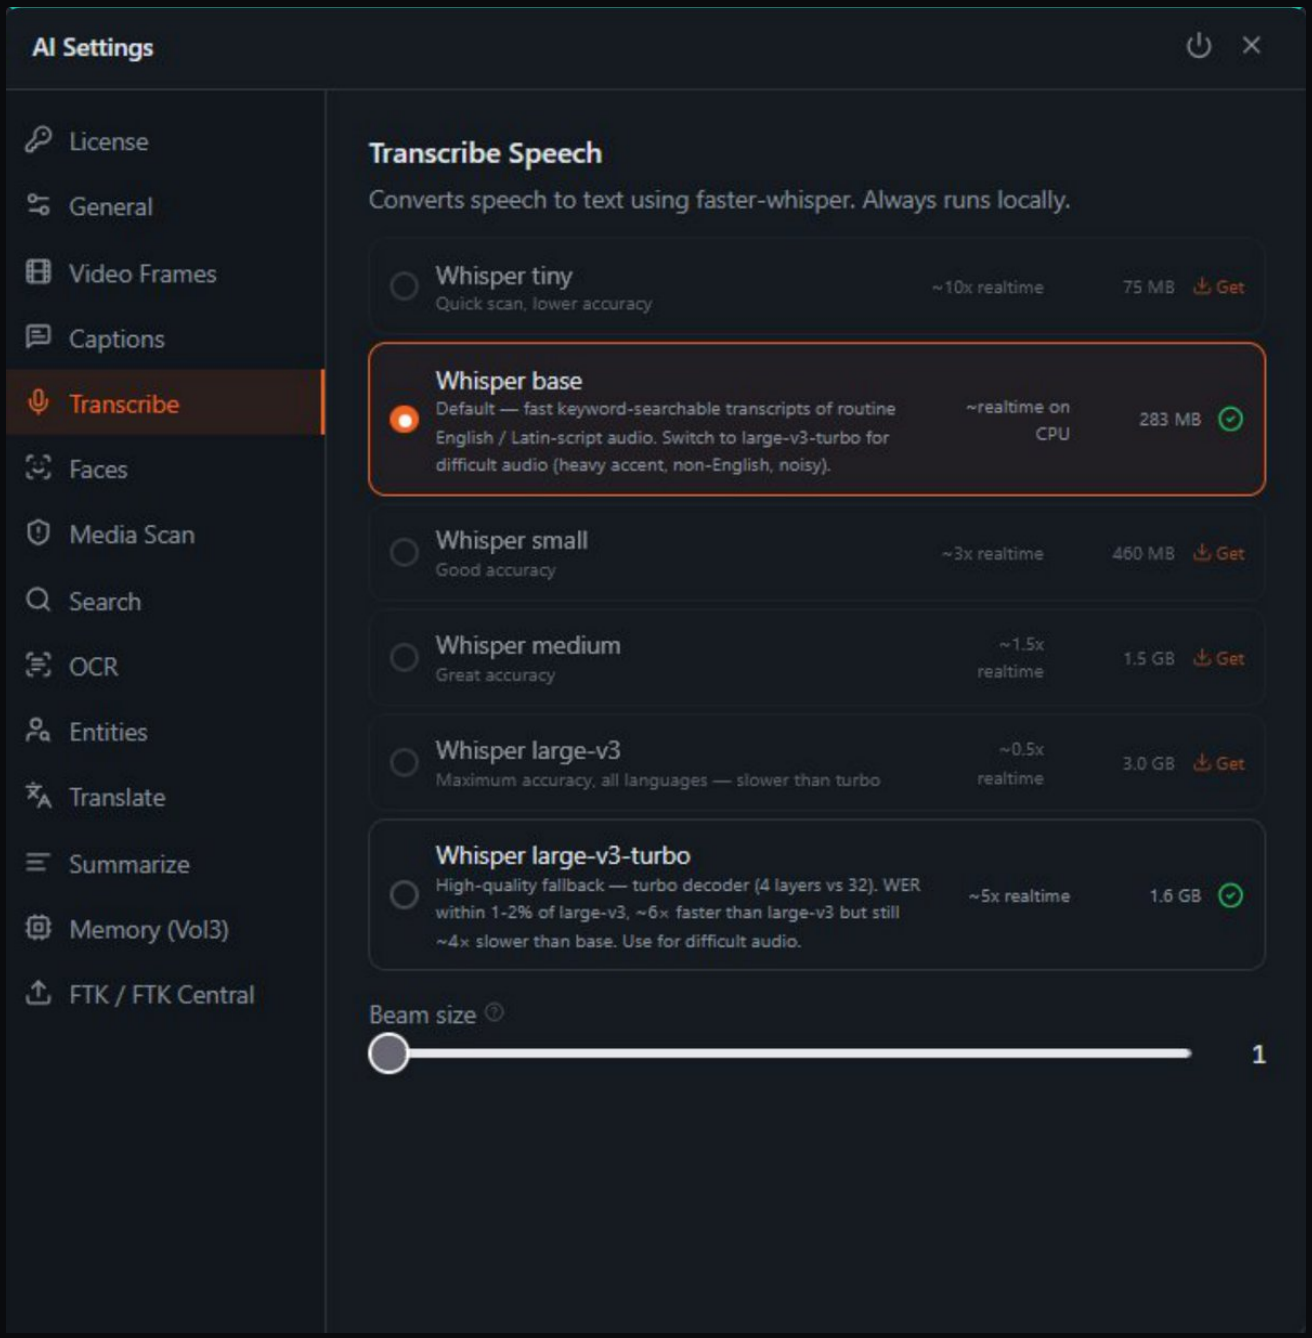

Whisper tiny Quick scan, lower accuracy. Useful for a rapid first pass when speed is the priority. Whisper base

Default

Fast, produces keyword-searchable transcripts for routine English and Latin-script audio. Recommended starting point. For difficult audio (heavy accent, non-English, noisy), switch to large-v3-turbo. Whisper small Good accuracy. A step up from base for cleaner transcripts without the full cost of larger models. Whisper medium Great accuracy.

Suitable when transcript quality is important and processing time is less of a concern. Whisper large-v3 Maximum accuracy across all languages. Slower than turbo. Use when the highest possible accuracy is required. Whisper large-v3-

turbo High-quality fallback using a turbo decoder (4 layers vs 32). Within 1-2% accuracy of large-v3 but ~6x faster. ~4x slower than base.

Best choice for difficult audio: heavy accents, non-English languages, or noisy recordings.

Beam Size

Controls the search width used during transcription decoding. Lower values are faster but may produce lower quality transcripts. Higher values are slower but can improve accuracy. This setting only has a meaningful effect on CPU. On GPU, transcription is fast at any beam size. Default: 1.

> For English evidence with clear audio, Whisper base is the right choice. For evidence with heavy accents, non-English speech, or background noise, switch to Whisper large-v3-turbo before running transcription.

Faces

Configures the face detection and clustering model. Always runs locally using InsightFace.

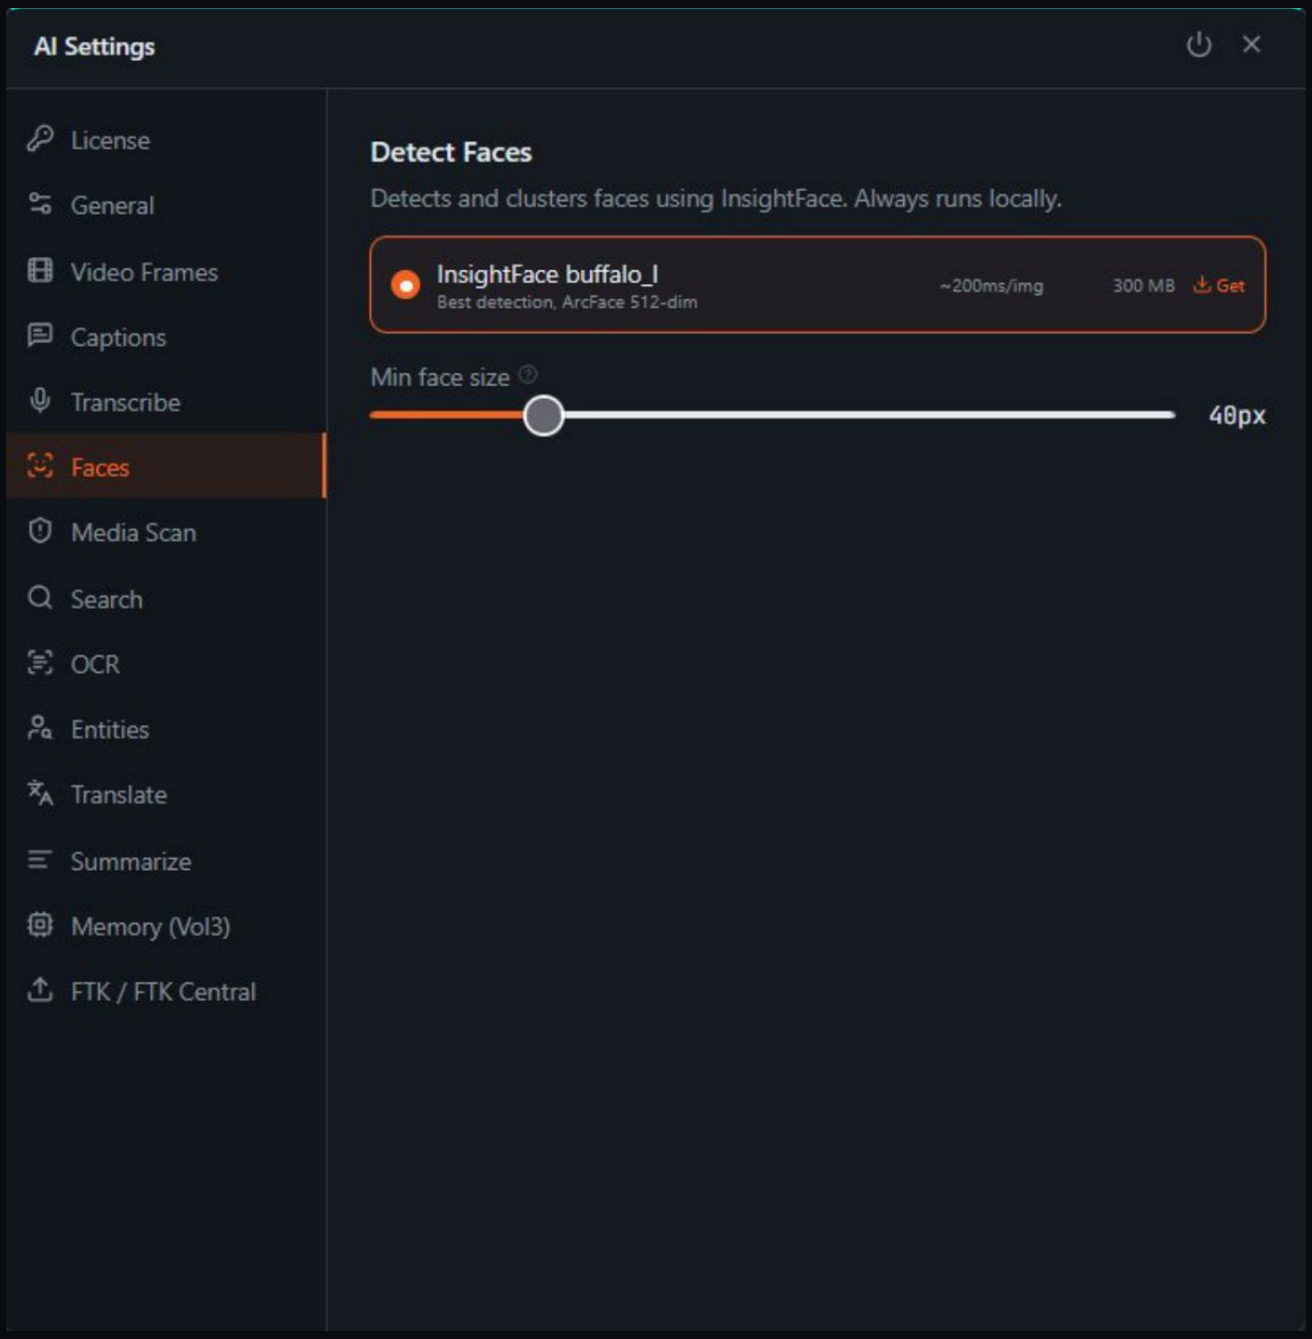

InsightFace buffalo_l Best detection model. Uses ArcFace 512-dimensional embeddings for accurate face clustering. Must be downloaded before face detection can be run.

Min Face Size

Sets the minimum pixel dimension for a face to be detected. Default: 40px. Lowering this value will catch smaller or more distant faces in images but may increase false positives. Increase it to reduce false positives when only clearly visible faces are relevant to the investigation.

Media Scan

Configures the SigLIP model used for visual image classification and manages the detection categories that appear in the Content filter. The first scan on a job locks the model. Later scans and visual search on that job reuse it.

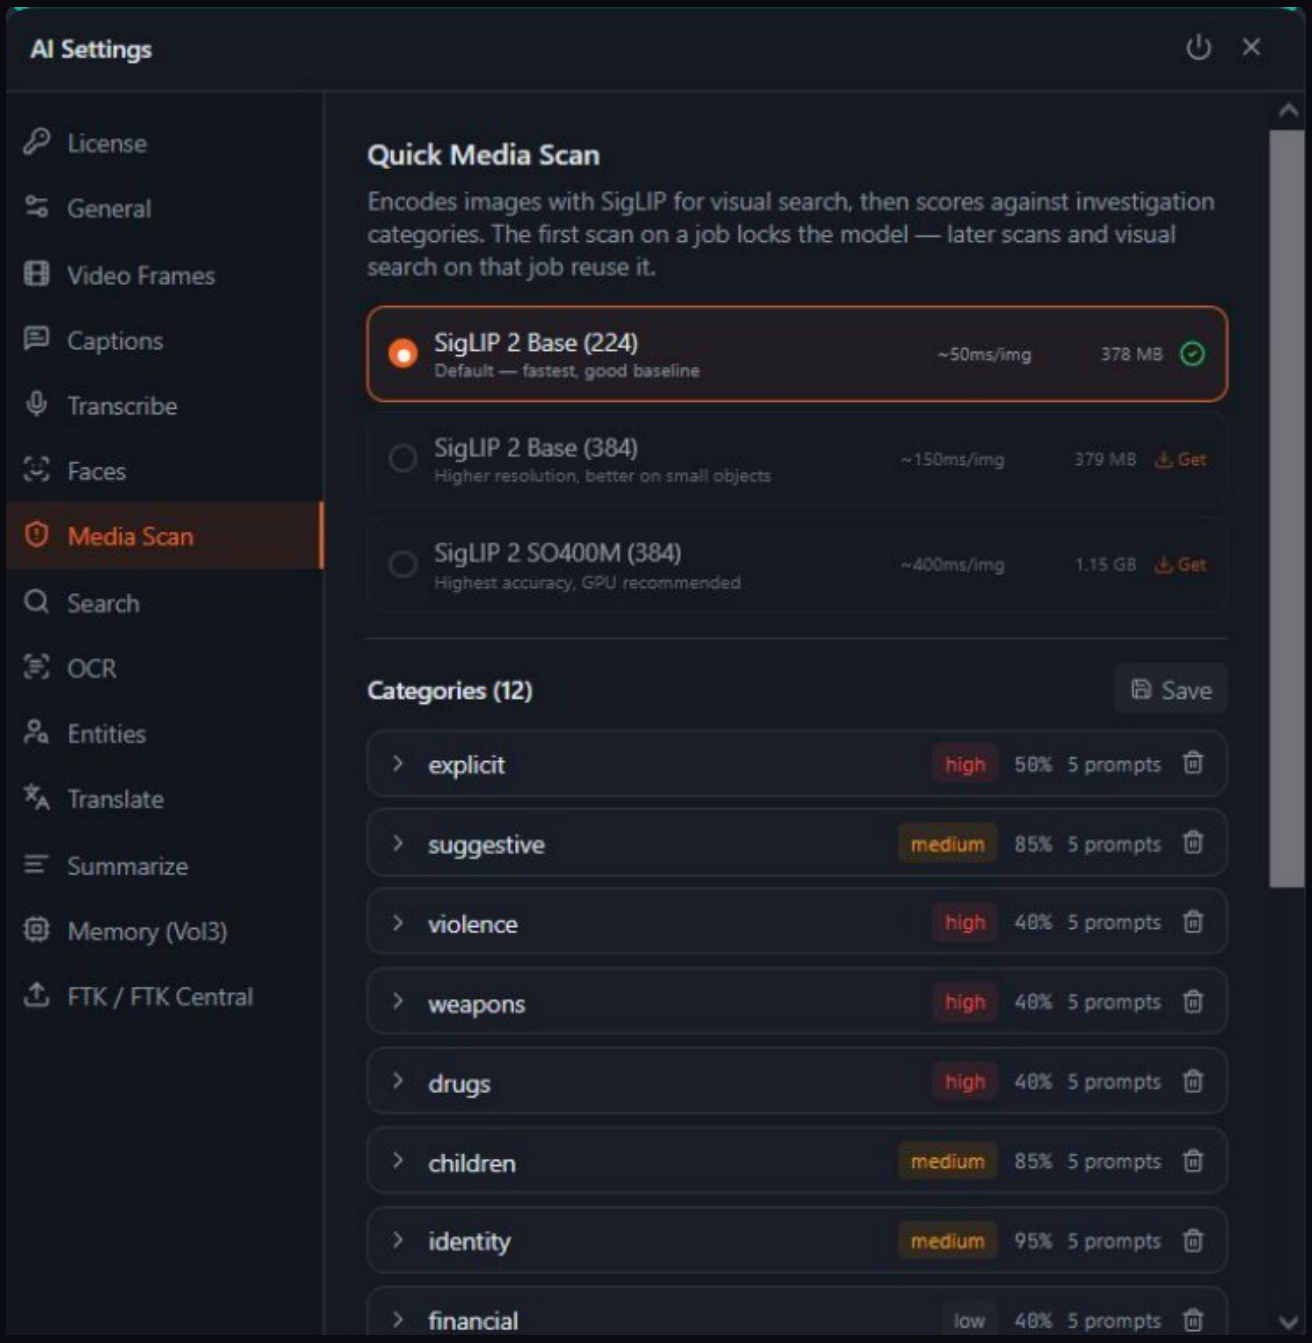

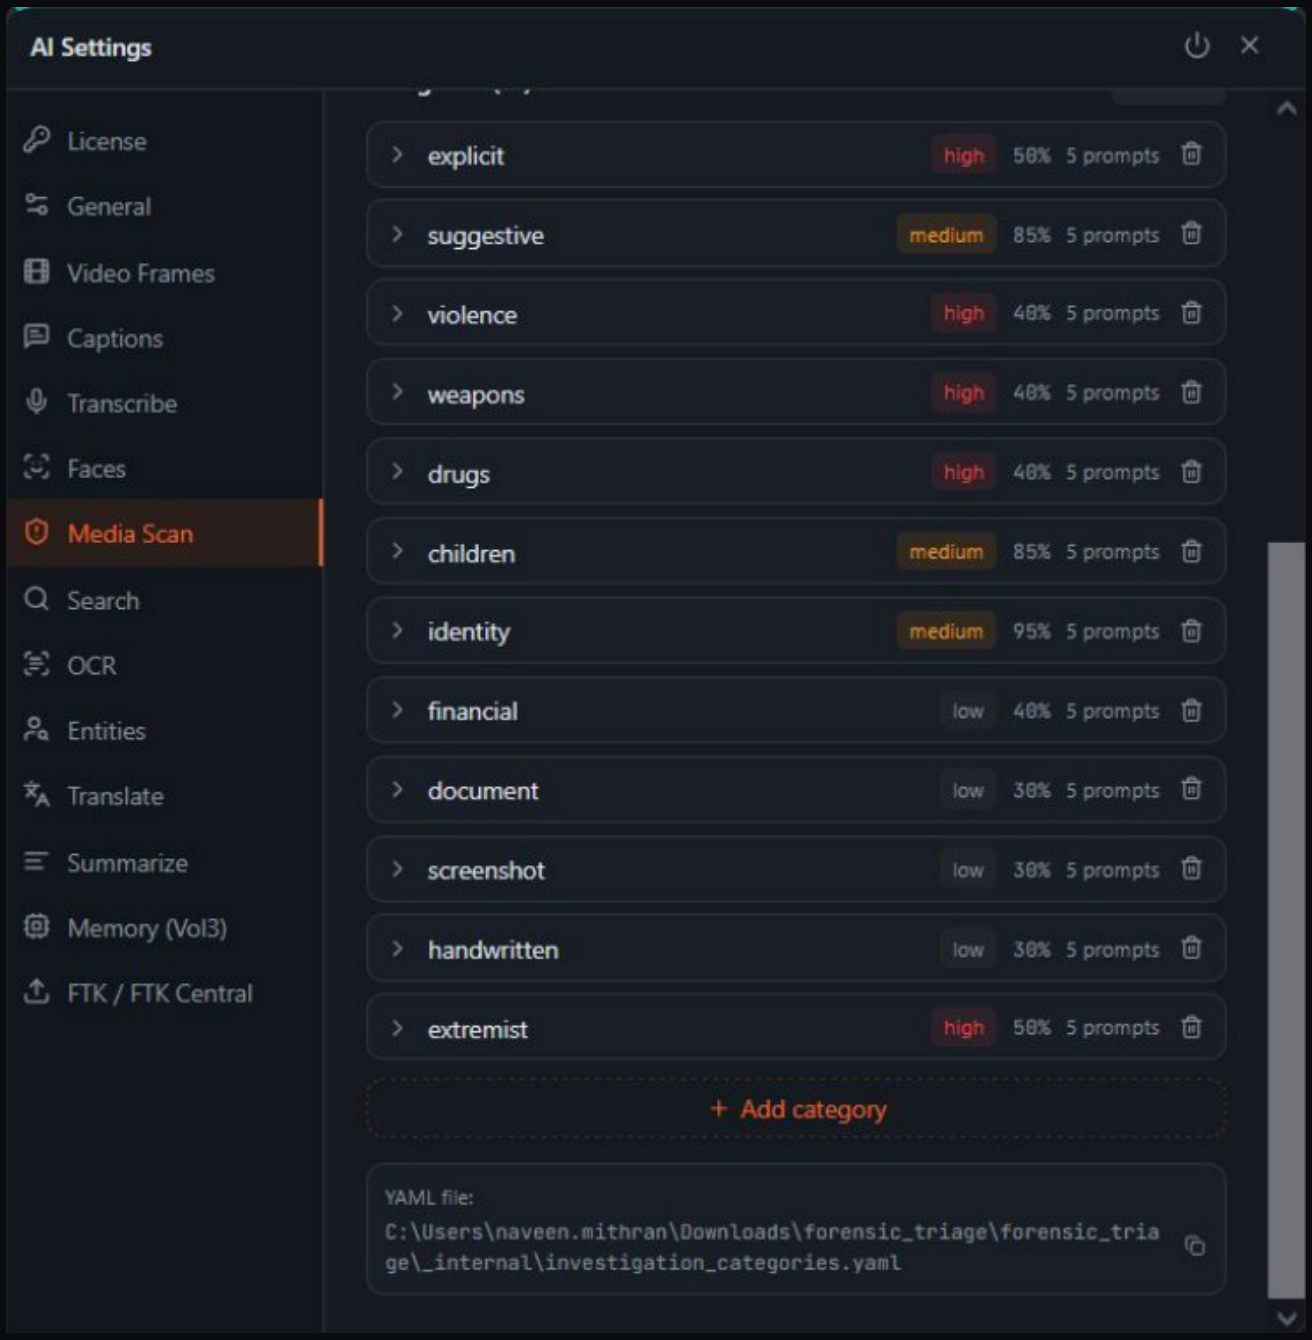

SigLIP 2 Base (224)

Default

Fastest, good baseline. Suitable for most investigations. SigLIP 2 Base (384) Higher resolution input, better at detecting content in small objects within images. SigLIP 2 SO400M (384) Highest accuracy. GPU recommended for practical performance.

Detection Categories

The 12 built-in categories are each configured with a severity level (high, medium, or low), a confidence threshold, and a set of visual prompts.

Each category shows the number of prompts used to detect it. explicit high severity, 50% threshold suggestive medium severity, 85% threshold violence high severity, 40% threshold weapons high severity, 40% threshold drugs high severity, 40% threshold children medium severity, 85% threshold identity medium severity, 95% threshold financial low severity,

40% threshold document low severity, 30% threshold screenshot low severity, 30% threshold handwritten low severity, 30% threshold extremist high severity, 50% threshold

Adding Custom Categories

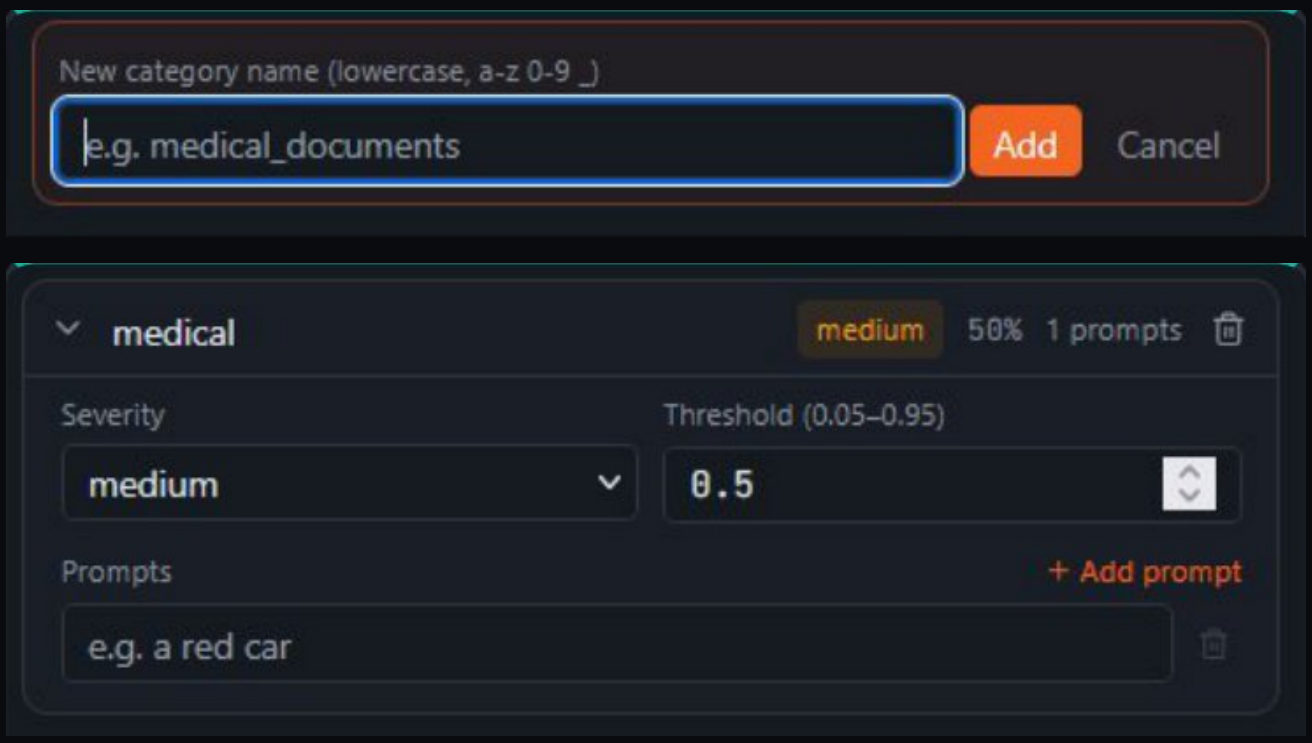

Click + Add category to create a custom detection category. Category names must be lowercase and can only contain letters, numbers, and underscores (e.g. medical_documents). Once created, each custom category exposes three configuration fields: Severity Set to low, medium, or high. Controls how the category is prioritised in results.

Threshold (0.05 to 0.95) The minimum confidence score required for a match to be reported. Default: 0.5. Prompts One or more plain-language descriptions of what the category looks like visually (e.g. "a red car"). Click + Add prompt to add multiple prompts. More prompts give the model more reference points for detection.

> All category configurations are stored in a YAML file in the ARMOURop app directory. Advanced users can edit this file directly to bulk-configure categories or share configurations across machines.

Search Settings

Configures the embedding model used for natural language search across captions, transcripts, and indexed document text. Also sets the confidence thresholds for semantic and visual search results.

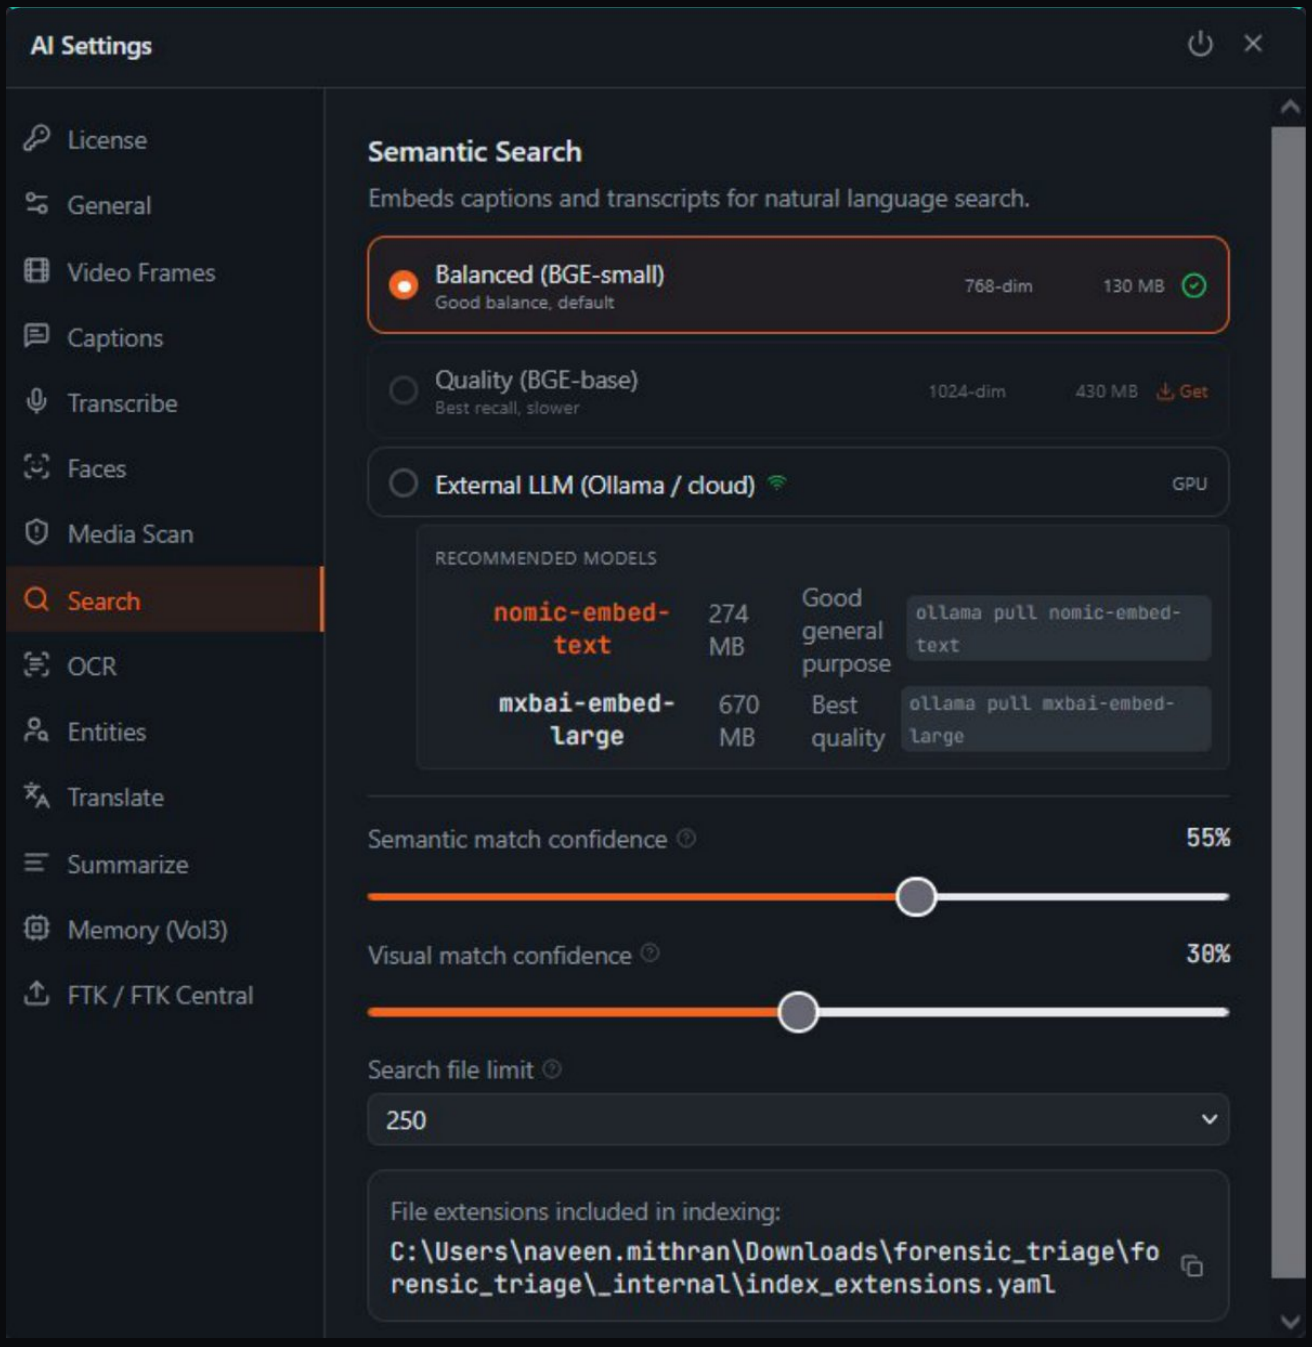

Balanced (BGE-small)

Default

Good balance between speed and recall. Downloaded and ready. Suitable for most investigations. Quality (BGE-base) Best recall, slower. Produces higher-dimensional embeddings for more accurate semantic matching. External LLM (Ollama / cloud) Use an Ollama-served or cloud embedding model.

Recommended options: nomic-embed-text (274MB, good general purpose) or mxbai-embed-large (670MB, best quality). Requires the Ollama endpoint to be configured in General settings.

Confidence Thresholds

Semantic match confidence Minimum confidence score for AI-driven semantic search results to be shown. Default: 55%. Lowering this shows more results but may include weaker matches. Raising it shows only high-confidence matches. Visual match confidence Minimum confidence score for SigLIP-based visual search results. Default: 30%.

Search File Limit

Sets the maximum number of files returned in a single search. Default: 250. Increase this if you want to see all results in a large evidence set, keeping in mind that higher limits may slow down search response.

File Extensions Included In Indexing

The list of file extensions that are indexed for document search is configured via a YAML file stored in the ARMOURop app directory. This determines which document types are included when Index Documents is run.

OCR

Configures the OCR model used to extract visible text from images and scanned documents. Uses RapidOCR running on CPU, supporting 80+ languages. Always runs locally.

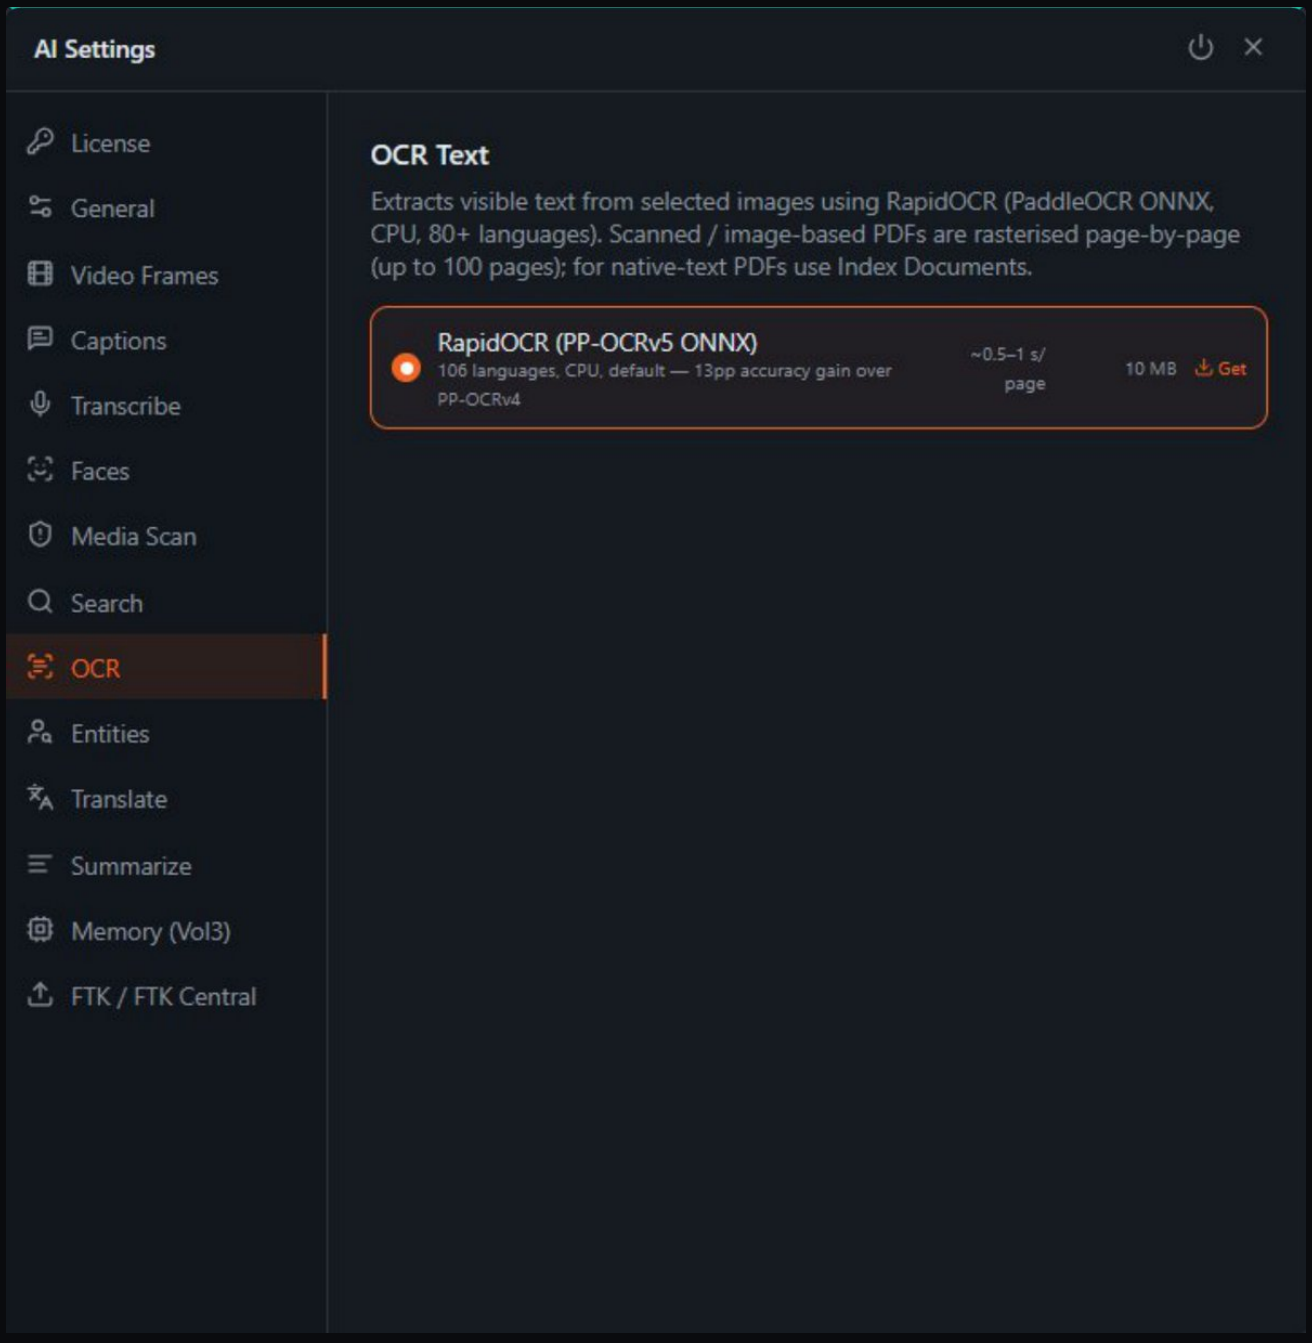

RapidOCR (PP-OCRv5 ONNX) The only available OCR model. Supports 106 languages on CPU. PP-OCRv5 delivers a 13 percentage point accuracy improvement over PP-OCRv4. Must be downloaded before OCR can be run.

> OCR is designed for image-based content: photos of documents, labels, signs, and ID cards. For native-text PDFs and Word documents where the text is already embedded digitally, use Index Documents instead. For scanned or image-based PDFs, OCR rasterises each page and reads it. Up to 100 pages per PDF are processed.

Entities

Configures the model used to extract forensically relevant named entities from chats, emails, documents, and OCR output. Detected entity types include person names, phone numbers, email addresses, IBANs, crypto wallet addresses, and vehicle plates.

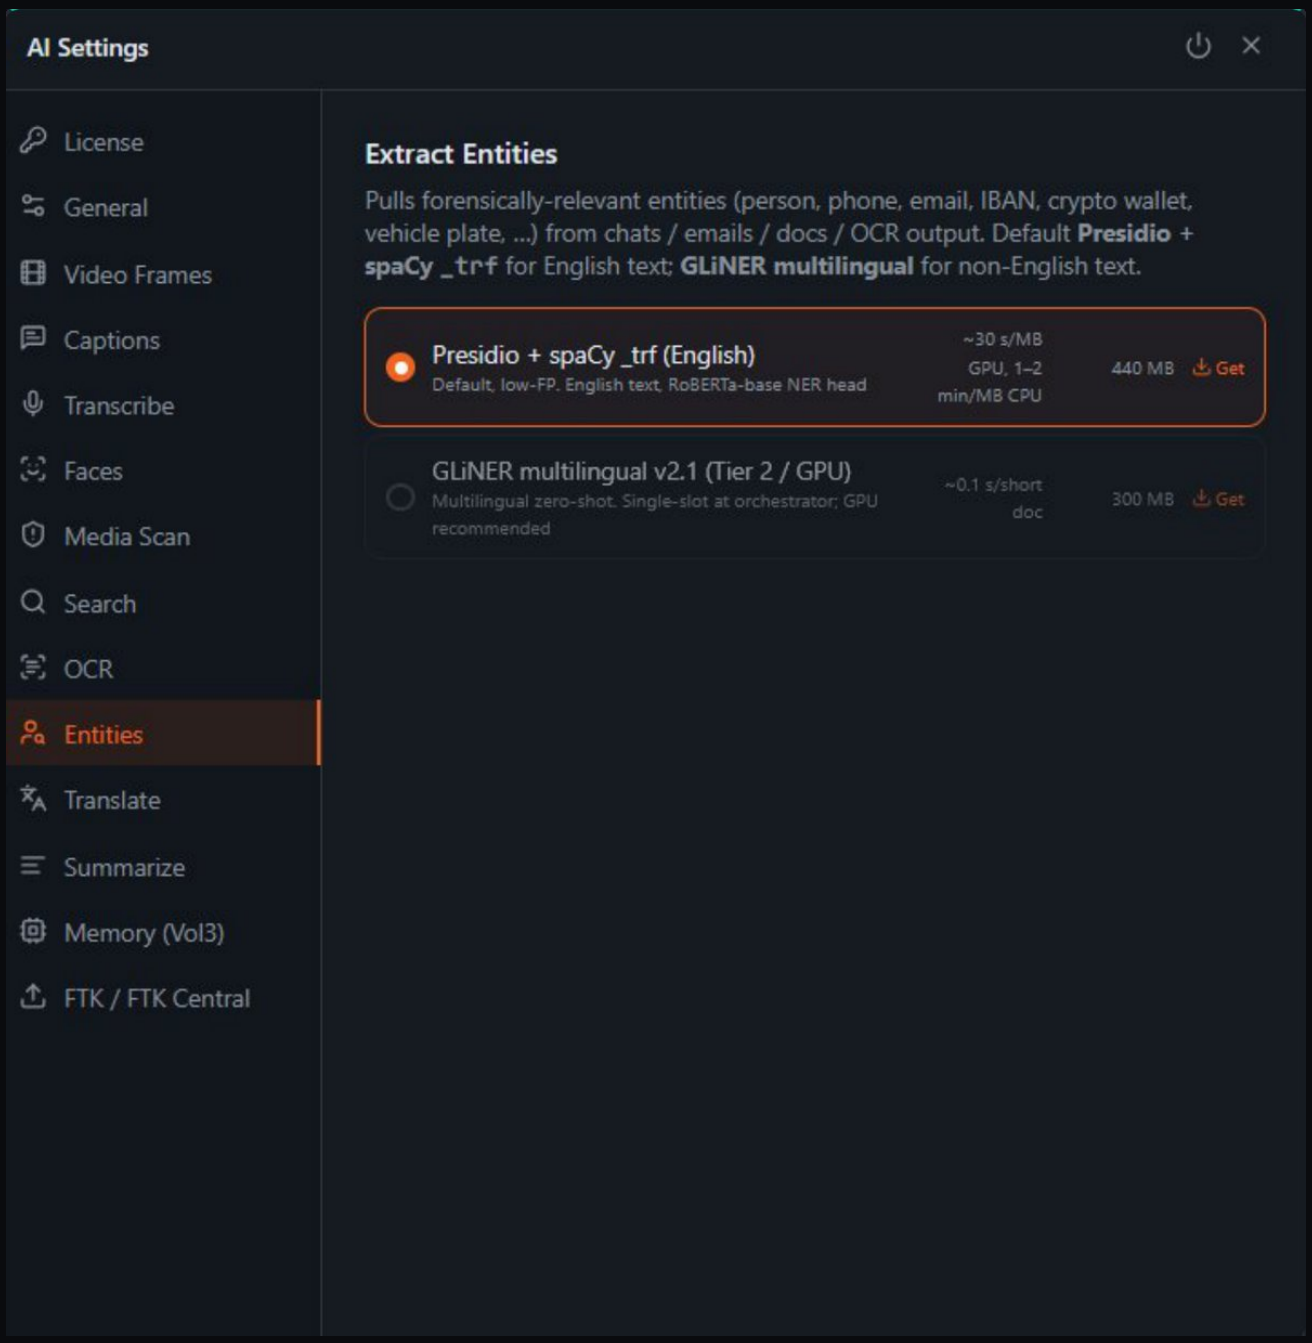

Presidio + spaCy _trf (English)

Default

Default for English text. Low false positive rate. Uses a RoBERTa-base NER head for accurate entity recognition in English-language content. GLiNER multilingual v2.1 (Tier 2 / GPU) Multilingual zero-shot model. Used automatically for non-English text when running alongside Presidio. GPU recommended. Single-slot orchestrator architecture.

The default behaviour uses both models together: Presidio handles English text and GLiNER handles non- English text automatically. Both models need to be downloaded before entity extraction can run.

Translate

Configures the translation model used to convert foreign-language documents to English. Auto-detects the source language. Always runs locally - no content is sent to external services.

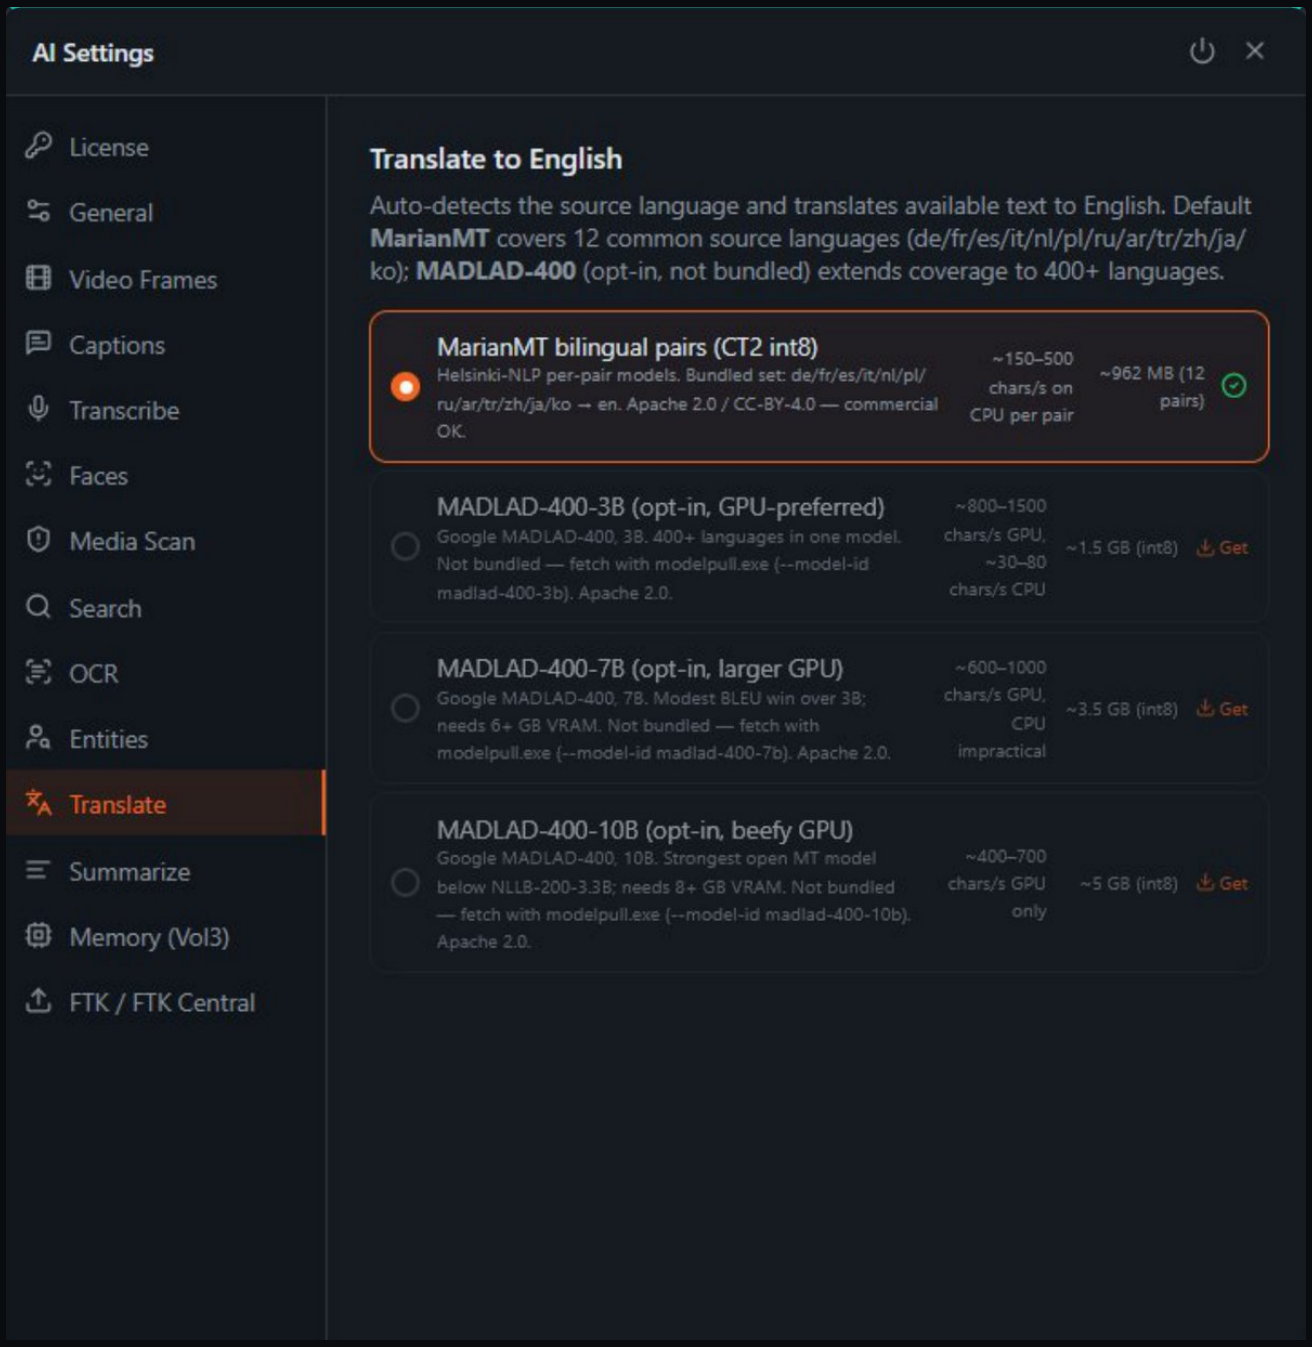

MarianMT bilingual pairs (CT2 int8)

Default

Bundled default covering 12 common source languages: German, French, Spanish, Italian, Dutch, Polish, Russian, Arabic, Turkish, Chinese, Japanese, and Korean. Licensed Apache 2.0 / CC-BY-4.0, commercial use permitted. MADLAD-400-3B (opt-in, GPU-preferred) Google MADLAD-400 3B model. Covers 400+ languages in a single model. GPU-preferred for practical speed.

Not bundled - fetch separately. Apache 2.0. MADLAD-400-7B (opt-in, larger GPU) Modest quality improvement over 3B. Requires 6+ GB VRAM. CPU use is impractical. Apache 2.0. MADLAD-400-10B (opt-in, beefy GPU) Strongest open translation model below NLLB-200-3.3B. Requires 8+ GB VRAM. GPU only. Apache 2.0.

> MarianMT is downloaded and ready for the 12 bundled languages. If your evidence contains languages outside that set, download MADLAD-400-3B for broad multilingual coverage.

Summarize

Configures the model used to generate AI summaries of documents and emails. Summaries are generated in batch mode - one summary per selected file.

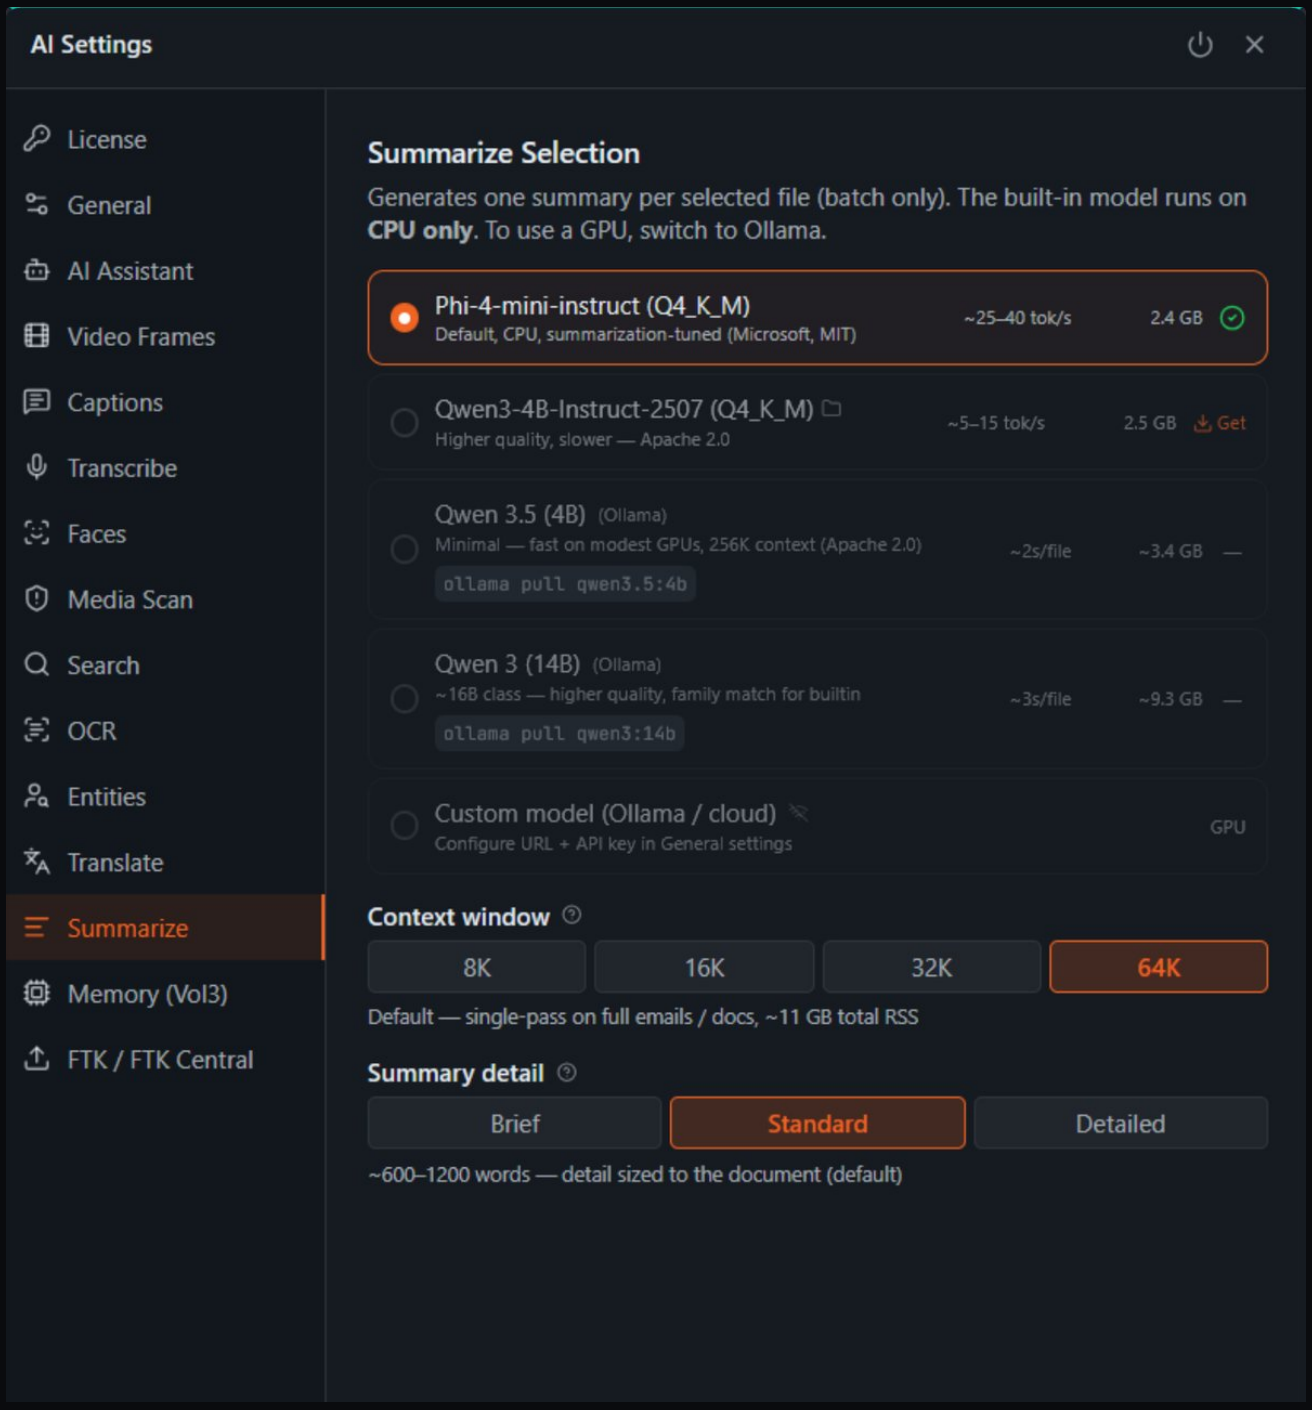

Phi-4-mini-instruct (Q4_K_M)

Default

Default, CPU-only, summarization-tuned model by Microsoft/MIT. ~25-40 tok/s, 2.4 GB, runs without Ollama. Qwen3-4B-Instruct-2507 (Q4_K_M) Higher quality summaries, slower than the default. ~5-15 tok/s, 2.5 GB. Apache 2.0. Qwen 3.5 (4B)

Ollama

Minimal footprint, fast on modest GPUs with a 256K context window. ~2s/file, ~3.4 GB. Apache 2.0. Install via ollama pull qwen3.5:4b . Qwen 3 (14B)

Ollama

~16B-class model, higher quality and a family match for the built-in Phi-4-mini default. ~3s/file, ~9.3 GB. Install via ollama pull qwen3:14b . Custom model (Ollama / cloud) Use any local Ollama tag or cloud model ID (e.g. gemini-2.5-flash, gpt-4o-mini). Requires GPU and an Ollama/cloud endpoint configured in General settings.

Context Window

Sets how much text the model processes in a single pass. Options: 8K, 16K, 32K, 64K. Default is 64K, which allows single-pass processing of full emails and documents at approximately 11GB total RSS memory usage. Reduce the context window if memory is constrained.

Summary Detail

Sets the target length of each generated summary. Options: Brief, Standard, Detailed. Default is Standard, which produces roughly 600-1200 words sized to the length of the source document.

Memory (Vol3)

Configures the YARA rules file used during memory analysis. YARA rules are applied by the vadyarascan and yarascan plugins when analysing memory dumps through Volatility 3.

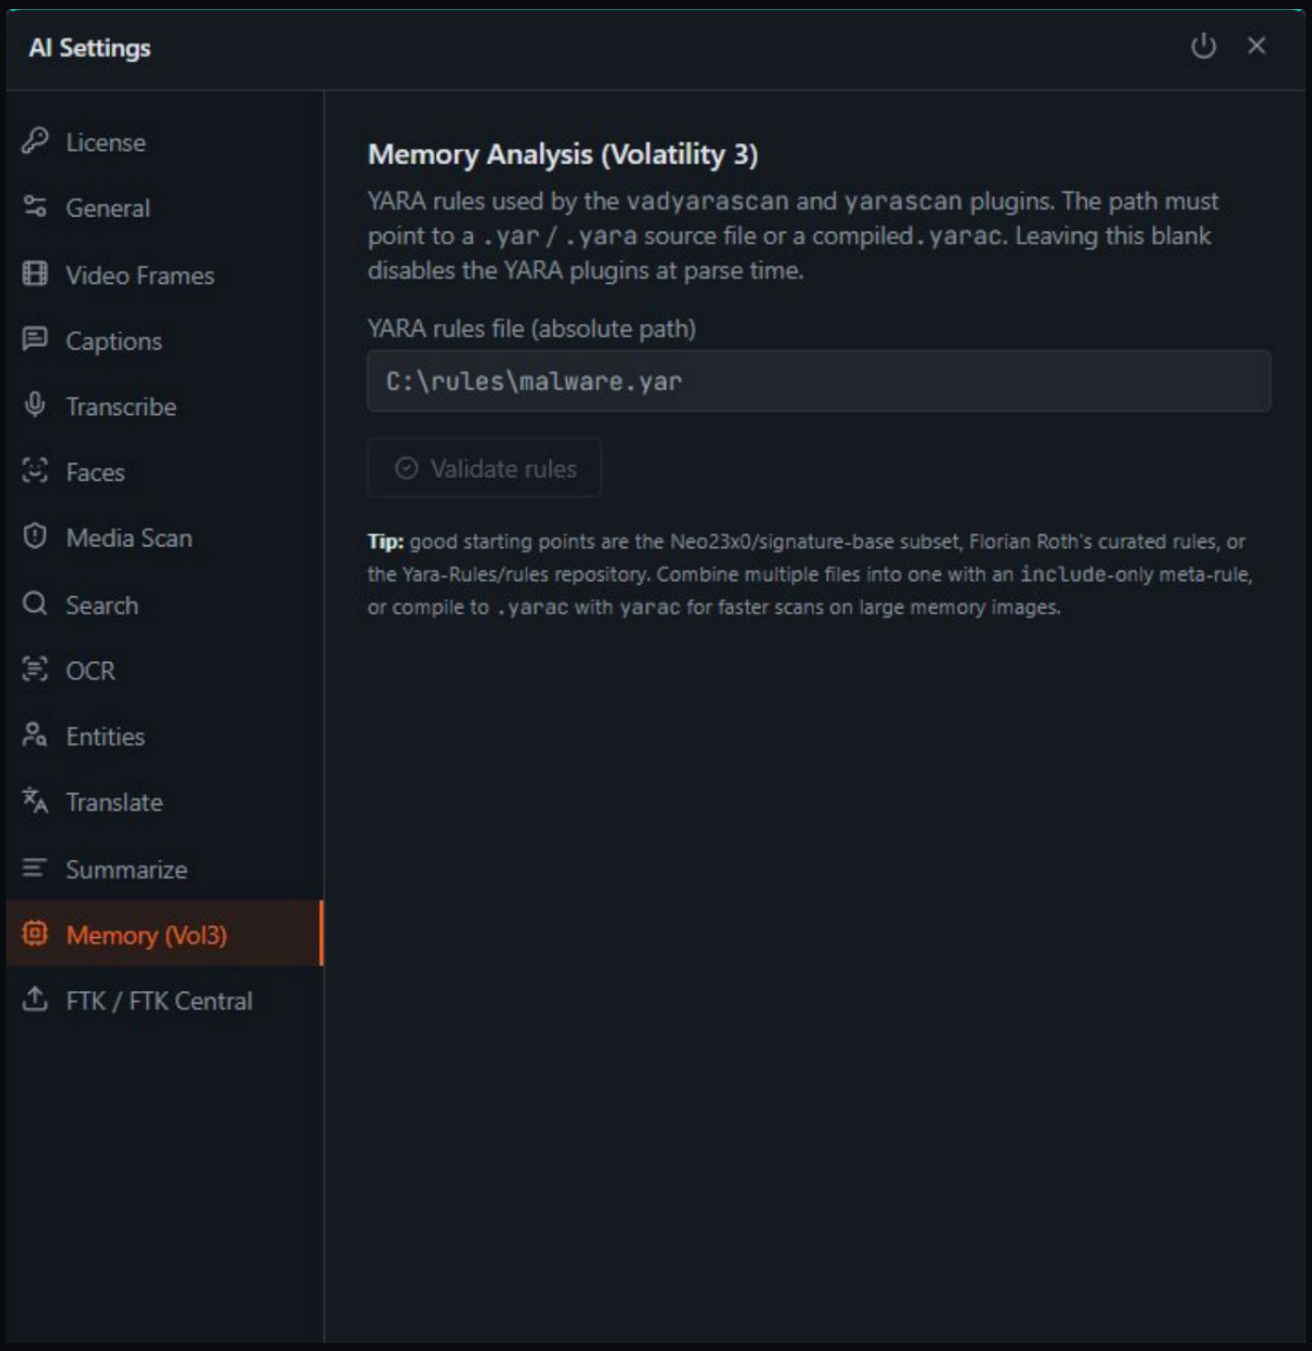

Yara Rules File

Enter the absolute path to a YARA rules file. The file can be a .yar or .yara source file, or a precompiled .yarac file. Leaving this field blank disables the YARA plugins at memory parse time.

1. Enter the rules file path Type or paste the absolute path to your YARA rules file (e.g. C:\rules\malware.yar).

2. Click Validate rules Click the Validate rules button to check that the file is valid before running a memory analysis job. This catches syntax errors before they affect analysis.

Good starting points for YARA rule sets include the Neo23×0/signature-base subset, Florian Roth's curated rules, or the Yara-Rules/rules repository. To use multiple rule files, combine them into a single file using an include-only meta-rule, or compile them to a .yarac file using yarac for faster scans on large memory images.

FTK / FTK Central

Configures the connection to an FTK or FTK Central server. This connection is required before bookmarked files can be exported from ARMOURop into an existing FTK or FTK Central case.

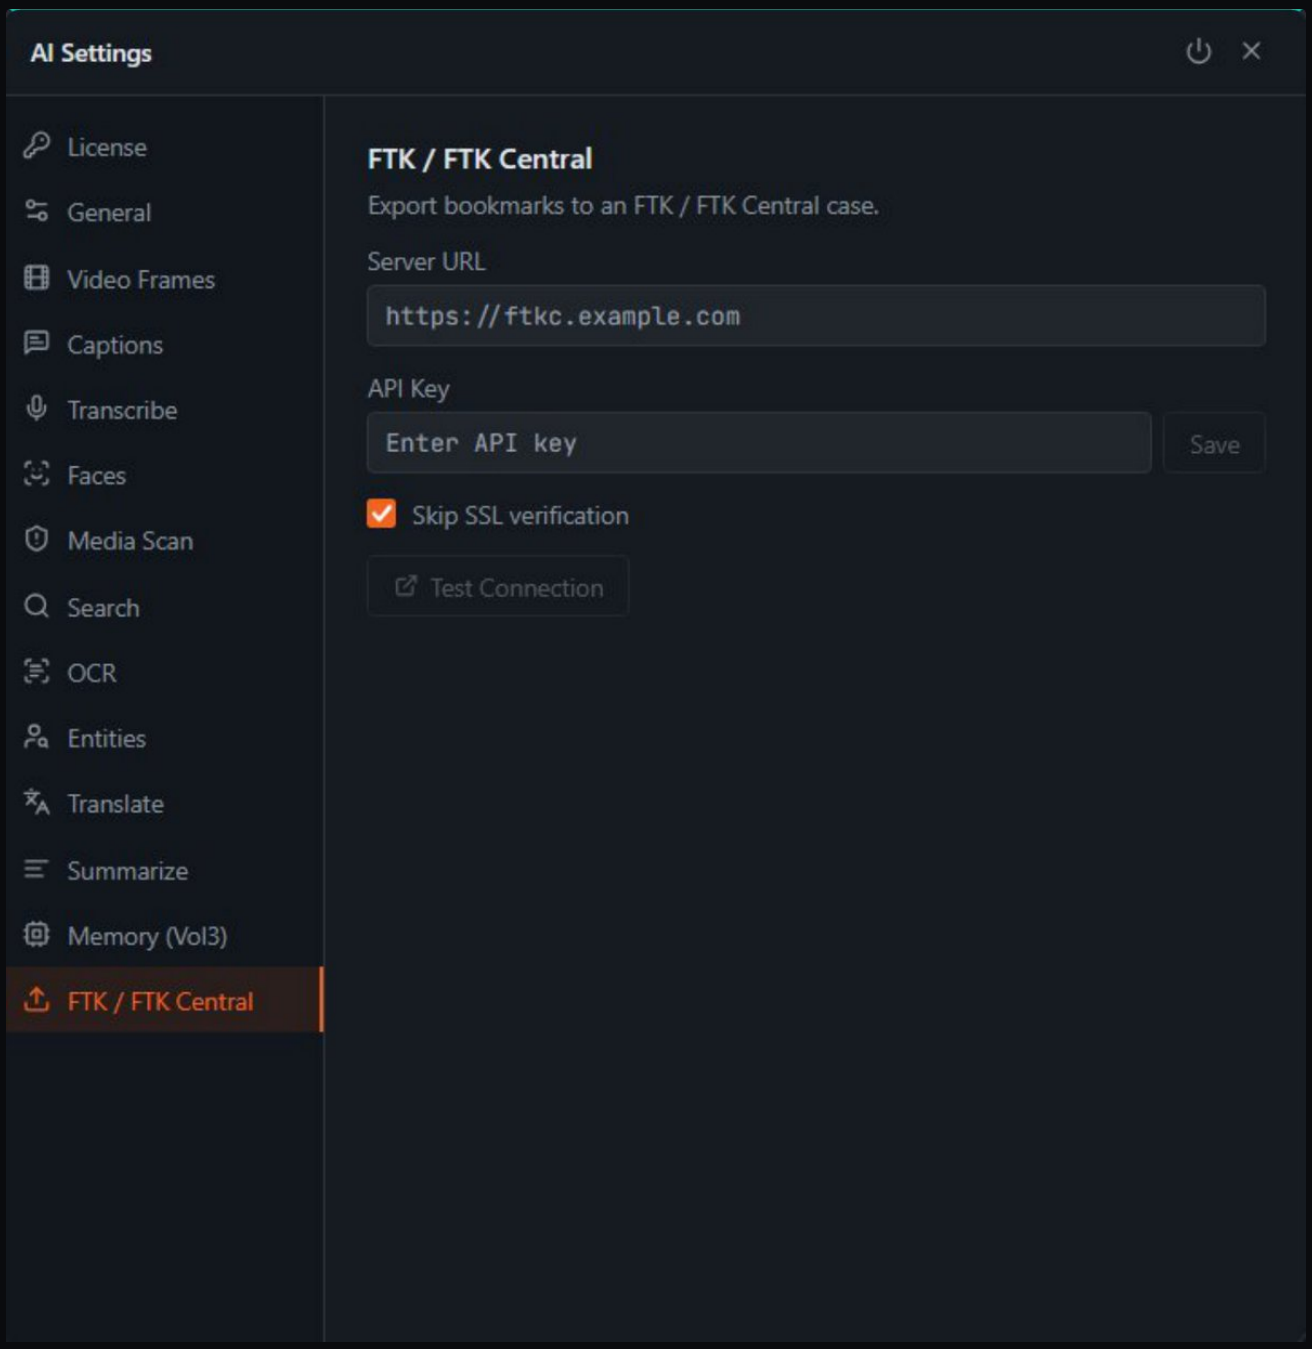

1. Enter the Server URL Enter the full URL of your FTK or FTK Central server (e.g. https://ftkc.example.com).

2. Enter the API Key Paste the API key for your FTK or FTK Central instance and click Save.

3. Skip SSL verification (if needed) Check this option if your FTK server uses a self-signed certificate or an internal CA that is not trusted by the machine running ARMOURop.

4. Test the connection Click Test Connection to verify that ARMOURop can reach the server with the provided credentials before attempting an export.

Once the connection is configured and tested, you can export bookmarked files to an FTK case from the Export menu in the Triage Console. See the Export to FTK / FTK Central section for the full workflow.