

Ariadne

Desktop forensic and operational intelligence workstation for evidence ingestion, external integrations, entity graphs, geospatial analysis, operational intelligence, eDiscovery, predictive research, AI-assisted analysis, and reporting.

User Manual

22.08.2026



Important notes and limitations

Ariadne accelerates triage and investigative navigation, but it does not replace human verification. All extracted entities, relations, dates, locations, face matches, and “contradictions” are heuristic outputs that must be reviewed against the original evidence.

The built-in Evidence viewer currently renders PDF pages only. Non-PDF files can be ingested (text extraction), but opening a page as evidence may fail unless the source is a PDF.

Operational Globe data can be delayed, incomplete, stale, geographically limited, or temporarily unavailable because several layers depend on external providers. Aircraft, vessel, satellite, weather, infrastructure, and imagery layers are analytical aids and are not certified navigation, traffic-control, collision-avoidance, or safety systems.

Correlation, proximity, encounter, pattern, anomaly, watchlist, replay, and orbital-opportunity outputs are derived analytical results. Spatial or temporal proximity does not establish causation, coordination, destination, affiliation, ownership, intent, or threat.

AI Analyst output is analytical assistance, not observed fact. When OpenAI is selected, the operational context chosen by the user is sent to the configured OpenAI service and may incur API charges. Review cited ARIADNE source records before relying on any AI-generated statement.

Entity Resolution (merge/split) and Ontology migrations are data-changing actions. Back up the database before applying them in a live investigation.

The External Ingestion API accepts material from authorised external integrations. Treat every API code as a sensitive credential, issue a separate code to each integration, and revoke it when it is no longer required. The operator remains responsible for lawful collection, access authority, privacy, retention, and source provenance.

PREDICTIVE output is probabilistic analytical assistance, not a fact, guarantee, diagnosis, or instruction to act. Review the probability type, uncertainty, confidence, supporting and contradicting evidence, data gaps, historical sample, and knowledge cut-off before relying on a result.

eDiscovery may use external web sources and the configured OpenAI service, may incur API charges, and may return incomplete, conflicting, outdated, privileged, or restricted material. Verify primary sources and observe confidentiality, privilege, copyright, data-protection, and evidential obligations.



Table of contents

This document uses Word heading styles. In Microsoft Word, you can generate an automatic table of contents via References → Table of Contents.

- 1. Getting started
- 2. Building the database (Ingestion and Indexing)
- 3. Exploring the case and Operational Globe
- 4. Using the External Ingestion API
- 5. Using PREDICTIVE
- 6. Using eDiscovery
- 7. Data quality and schema management (Resolution and Ontology)
- 8. Exporting and reporting
- 9. Keyboard and mouse reference
- 10. Troubleshooting
- Appendices



1. Getting started

1.1 What Ariadne does

- Ingest evidence from the **files/** folder (PDFs and common office/text formats).
- Extract text (including OCR for image-based PDFs when Tesseract is available).
- Extract entities (people, organizations, locations, events, accounts) and build a relationship graph.
- Extract identifiers (email, URL, domain, IPv4, IBAN, Portuguese NIF, phone, hashes) and link them to entities.
- Detect faces in PDF pages (optional) and provide a Face Gallery with thumbnails per person.
- Generate investigation views: Events Timeline, Geo Map, Heat Map, Contradiction Radar, Alerts and Patterns.
- Produce exports: alerts (CSV/XLSX), dossiers (Markdown/HTML/CSV), intelligence reports (DOCX), operational briefings (HTML/JSON), AI analyses (HTML), and AOI geometry (GeoJSON/KML).
- Use the interactive Operational Globe for satellite, aircraft, vessel, weather, environmental, infrastructure, imagery, correlation, replay, and encounter analysis.
- Create advanced Areas of Interest (AOIs) using circles, rectangles, polygons, and corridors, with local geofencing and GeoJSON/KML import/export.
- Analyze local operational history using correlation, pattern/anomaly detection, watchlists and alerts, replay, briefings, and encounter/proximity analysis.
- Use AI Analyst to ask natural-language questions or generate reports from the operational data you select.
- Receive authorised camera, image, PDF, e-mail, and WhatsApp material through the External Ingestion API.
- Use PREDICTIVE to investigate future-oriented questions, compare evidence, run backtests and what-if scenarios, and review probability history.
- Use eDiscovery to conduct structured external research, preserve relevant sources, maintain Living Dossiers, test claims, and prepare evidence products.

1.2 Typical investigative workflow

1. Place evidence into the correct subfolders under files/ (PDFs, Emails, Finance, etc.).
2. Open Settings, choose the AI mode (Local or OpenAI), and configure the selected AI and any operational-source credentials you intend to use.
3. Run Ingestion and Indexing -> Build Database (first run) or Reindex New/Changed Only (updates).



4. Optionally start the External Ingestion API and connect authorised external programs using a separate API code for each integration.
5. Use Relationships, Smart Search, Events Timeline, Geo Map, Heat Map, Contradiction Radar, and Faces Intelligence to explore the investigative corpus.
6. Use PREDICTIVE for evidence-based forecasts, historical cut-off questions, backtesting, what-if analysis, anomaly review, and prediction history.
7. Use eDiscovery when the task requires structured external research, source preservation, a Living Dossier, claim testing, or a formal evidence product.
8. Open Operational Globe, define or load an Active AOI, then enable only the intelligence layers needed for the current task.
9. Refresh relevant operational sources such as Aircraft, Vessels, Weather, Infrastructure, Satellite/Imagery, then use Correlation, Patterns, Watchlists, Replay, Briefing, and Encounters as needed.
10. Use AI Analyst after loading the operational data relevant to your question. Review the selected context, click Analyze, then verify important statements against the records shown in SOURCES.
11. Optionally clean up identity duplication using Resolution and keep the investigative schema consistent using Ontology.
12. Curate relevant investigative and operational observations into an Investigator Dossier.
13. Export the required Dossier, Intelligence Report, PREDICTIVE evidence, eDiscovery product, Operational Briefing, AI Analysis, AOI geometry, or other supported output for reporting and handoff.

1.3 System requirements

- Windows 10/11 (portable folder deployment).
- CPU: modern 4-core+ recommended; RAM: 16 GB+ recommended for large corpora.
- Disk: allow space for the database and face thumbnail cache (can grow with large PDF sets).
- Optional: Tesseract OCR installed for scanned PDFs (otherwise OCR is skipped).
- Optional: OpenAI API key if using OpenAI mode or AI Analyst with OpenAI. PREDICTIVE Deep Search and eDiscovery require a valid OpenAI configuration and supported model access.



- Network access is required for live/external operational data and for OpenAI mode. Local investigation features and previously available local data remain usable offline where supported.
- Optional provider credentials are required for commercial or authenticated sources such as Datalastic, Planet, Maxar/Vantor, and ICEYE. AISStream.io also requires its own API key.
- For large Operational Globe workloads, a modern GPU and stable network connection are recommended but not mandatory for the local engine.

1.4 Folder layout (inputs)

Ariadne expects a folder named files/ next to the application executable.

Path (relative to application folder)	Purpose / notes
files/PDFs/	PDF evidence. Best supported format: evidence viewer, page references, faces, OCR.
files/Logs/	System logs, exports, CSV/JSON logs. Useful for timeline/log extraction.
files/CRM/	Contacts, CRM exports, notes.
files/ERP/	ERP exports (orders, invoices, inventory).
files/Finance/	Statements, payments, invoices, bank evidence.
files/WhatsApp/	Chat exports.
files/Email/	Email exports.
files/OSINT/	Web captures, OSINT reports, intelligence notes.
files/_tmp/	Ignored by ingestion (use for staging).
files/_ignore/	Ignored by ingestion (use for non-case materials).
files/ (root)	Legacy: PDFs placed directly under files/ are still ingested.

1.5 Supported file formats

Ariadne ingests PDFs and several common text/office formats. Non-PDF files are treated as a single text page (no rendering in the Evidence viewer).

Format	Ingestion	Notes
PDF	Full: page text + optional OCR + faces + page evidence	Recommended for primary evidence.
TXT / MD	Text extraction	Single “page”.
CSV / JSON / XML	Text extraction	Single “page”. Large files



		may increase ingestion time.
HTML	Text extraction	Basic tag stripping.
DOCX	Text extraction	Extracts paragraph text.
XLSX	Text extraction	Extracts cell text across sheets.

1.6 Where Ariadne stores its data

By default, Ariadne stores state under your Desktop:

Path	Contents
Desktop/ariadne-bd/ariadne-storage/ariadne.db	Main SQLite database (entities, relations, passages, events, alerts, etc.).
Desktop/ariadne-bd/ariadne-storage/settings.json	UI and runtime settings.
Desktop/ariadne-bd/ariadne-output/	Logs and exports (alerts CSV/XLSX, dossiers, intelligence report DOCX, etc.).
Desktop/ariadne-faces/	Face cache DB (faces.db) and optional face crops/thumbnails.
Desktop/ariadne-bd/ariadne-storage/operational-cache/	Saved AOIs, watchlists, operational history, and cached operational data used by supported workspaces.
Desktop/ariadne-bd/ariadne-storage/evidence/satellite-imagery/	Locally archived satellite-imagery previews or products created through supported imagery workflows.

1.7 Selecting AI mode (Local vs OpenAI)

Ariadne can run with a local model or with OpenAI models. The selected AI configuration is shared by AI-assisted features, including AI Analyst and standard PREDICTIVE research.

Mode / setting	When to use	Trade-offs / notes
Local AI	Air-gapped or privacy-sensitive workflows; no cloud AI request for Local AI tasks.	Quality and speed depend on the installed local model and hardware.
OpenAI AI	Cloud-assisted extraction, Smart Search, Intelligence Report, and AI Analyst when the selected OpenAI model is available to the configured account.	API usage is billable and depends on model, context size, reasoning effort, output length, and provider pricing.
OpenAI model	Choose the OpenAI model available in Settings and permitted for your OpenAI account.	Model availability, speed, capability, and pricing can change. Test the selected AI before a large analysis.



Reasoning effort	Use a higher reasoning level for tasks that genuinely need deeper analysis.	Higher reasoning can increase latency and OpenAI usage cost. Start lower while testing a workflow.
-------------------------	---	--

When OpenAI is selected, configure the API key, model, and reasoning effort in Settings, then click Save. Use Test Selected AI in AI Analyst to confirm the configuration before running an analysis.

PREDICTIVE Deep Search and eDiscovery use the OpenAI configuration indicated in their respective workspaces. Open each workspace and confirm that every configuration status is ready before starting a long-running operation.



2. Building the database (Ingestion and Indexing)

2.1 What “Build Database” does

- Scans files/ (all subfolders except _tmp and _ignore).
- For PDFs: reads per-page text; runs OCR on pages with no text if Tesseract is available; stores per-page snippets for citations.
- For non-PDF files: extracts text and stores it as a single “page”.
- Runs identifier extraction and links identifiers to entities.
- Runs AI extraction (Local/OpenAI) to propose entities, relations, and events.
- Optionally runs face detection/embedding on PDF pages and stores face occurrences.
- Updates indexes used by Smart Search and other tabs.

2.2 Controls in the Ingestion and Indexing tab

Buttons and options:

- Build Database — full rebuild pass over all input files.
- Reindex New/Changed Only — incremental update for new or modified files.
- Cancel — stops ingestion as soon as the current unit of work is safe to stop.
- Enable face search during ingestion — enables PDF face detection and thumbnail generation.
- Resume interrupted build from checkpoint — resumes from the last checkpoint after a crash/cancel.

Progress bars show file-level and page-level progress. The detailed log reports processing status, OCR activity, extraction progress, and any errors that require attention.

2.3 OCR for scanned PDFs

If a PDF page has little or no embedded text, Ariadne can run OCR using Tesseract. If Tesseract is not installed or not configured, OCR is skipped and that page may remain effectively “silent” for search.

Configuration options:

- Install Tesseract OCR on the machine.
- If Ariadne cannot find it, set the environment variable `ARIADNE_TESSERACT_CMD` to the full path of `tesseract.exe`.



2.4 Face detection during ingestion

When “Enable face search during ingestion” is enabled, Ariadne scans PDF page images for faces, generates embeddings, and stores face occurrences. Face thumbnails (crops) may be written under Desktop/ariadne-faces/.

If the system prompts for labels during ingestion, provide consistent naming (e.g., “John Smith”, “Unknown Male 01”). You can refine identities later using Resolution (merge/split).

2.5 Incremental updates and evidence hygiene

- Prefer immutable evidence: do not edit PDFs in place after ingestion. If you must, reindex afterwards.
- Renaming/moving files may cause them to be treated as new evidence; keep stable paths during an investigation.
- Use files/_tmp or files/_ignore for staging materials you do not want ingested.
- Back up Desktop/ariadne-bd/ariadne-storage/ariadne.db before major operations (merges, migrations).



3. Exploring the case and Operational Globe

3.1 Tabs at a glance

Tab	Primary purpose
Relationships	Interactive entity graph + evidence drill-down.
Smart Search	Ask questions; Ariadne answers with citations you can open as evidence.
Faces Intelligence	Known people list, occurrences, and Face Gallery thumbnails.
Events Timeline	Chronological timeline plus an interactive event graph view.
Geo Map	Map view of extracted locations with evidence links.
Heat Map	Heat/intensity map (People/Locations/Events) over the basemap.
Contradiction Radar	Heuristic detection of conflicting values/amounts/dates for the same identity.
Alerts and Patterns	Rule-based alerts, triage workflow (open/closed), exports.
Investigator Dossier	Curated case bundle; export Markdown/HTML/CSV.
Intelligence Report	Iterative AI narrative + exports to DOCX.
Settings	AI mode/key/model + diagnostics and paths.
Resolution	Merge/split entities to clean duplicates and conflations.
Ontology	Active ontology version + migrations + dataset stats + validation.
Operational Globe	Interactive operational workspace for AOIs, satellites, imagery, aircraft, vessels, weather, infrastructure, correlation, patterns, watchlists, replay, briefings, encounters, and AI Analyst.
API	Manage the External Ingestion API, authorised integrations, permissions, credentials, and recent ingestion activity.
PREDICTIVE	Ask future-oriented questions, investigate evidence iteratively, review probabilities, run backtests and what-if scenarios, and examine prediction history.
eDiscovery	Conduct structured external research, create and maintain evidence dossiers, test claims, compare dossiers, and generate evidence products.



3.2 Relationships

Use Relationships to explore how entities connect across the corpus. The graph is built from entity co-occurrence and extracted relations.

Typical use:

- Select a focus entity (person/organization/location/etc.).
- Choose the depth (how many hops away from the focus entity).
- Build/refresh the graph.
- Click nodes/edges to view supporting evidence. Use “Open Evidence” to render the source PDF page.

If the graph feels noisy, reduce depth and/or use Resolution to merge duplicates before expanding.

3.3 Smart Search

Smart Search is a question-answer interface over the ingested corpus. Ariadne retrieves relevant passages and uses the configured AI mode to answer, always returning citations (file + page) for auditability.

- Enter your question in the text box (English or Portuguese).
- Choose Language: auto / en / pt (controls the answer language).
- Click Ask.
- Review the Answer section, then select a row in Citations and click Open Evidence.

Caveat: Smart Search answers can be wrong or incomplete. Treat them as navigation aids; verify against cited pages.

3.4 Faces Intelligence

Faces Intelligence is available when face detection is enabled during ingestion and the evidence is PDF-based.

- Select a known person from the list to view face occurrences across evidence pages.
- Use Open Evidence to open the PDF page for a selected occurrence.
- Use Open Face Gallery to browse thumbnails for the selected person and jump to the evidence page behind each thumbnail.

If thumbnails are not available, it usually means face crops were not generated (face search disabled) or the source pages were not renderable.



3.5 Events Timeline

Events Timeline aggregates extracted timeline items (e.g., flight, check-in/check-out, contacts, generic events) and provides both a chronological list and an interactive graph view.

Filters and actions:

- Set From/To dates to constrain the timeline window.
- Filter by Entity (ALL or a specific identity).
- Filter by Type: ALL / flight / checkin / checkout / contact / event / other.
- Click Refresh to rebuild the view.
- Use Reset Graph View if panning/zooming leaves the graph off-screen.
- Select an event row and use Open Selected Evidence to open the supporting PDF page.

The timeline graph supports panning/zooming and is intended for fast navigation between clustered events. Always open and verify the underlying evidence.

3.6 Geo Map

Geo Map plots extracted locations over a world basemap and links them back to evidence. It relies on a basemap image (world.jpg/png/etc.) and a locations catalogue (world_locations.json).

Prerequisites:

- Basemap image: place world.jpg (or world.png/world.webp) next to the executable if it is not bundled.
- Locations catalogue: world_locations.json can be placed next to the executable, under files/geo/, under files/, or under resources/.
- Optional: set ARIADNE_GEO_CATALOG to override the catalogue path explicitly.

Using Geo Map:

- Choose Identity = ALL or a specific identity (limits markers to that identity's evidence).
- Choose Type = ALL / PERSON / ORGANIZATION / LOCATION (controls which extracted mentions are mapped).
- Set Min confidence and Min score to reduce noise.
- Optionally enable Use date range, then set From/To.
- Optionally enable Person↔Location links to draw lines between people and locations found together.
- Click Apply Filters or Refresh Geo Cache when inputs change significantly.



- Navigation: pan with middle mouse (or Alt+left drag); zoom with mouse wheel or +/-; use Fit to Data / Reset View as needed.

3.7 Heat Map

Heat Map visualizes intensity clusters across the corpus (People / Locations / Events), helpful for finding hotspots at a glance.

Using Heat Map:

- Select the layer: People, Locations, or Events.
- Set From/To date range (optional, depending on your corpus).
- Tune Min confidence and Min score to reduce noise.
- Click Refresh to recompute the heat layer.
- Navigation: click-drag to pan; zoom with mouse wheel or +/-; press 0 to reset zoom.
- Select a row in the evidence table (if shown) and open evidence to validate.

3.8 Contradiction Radar

Contradiction Radar scans extracted events/log-like facts and highlights potential contradictions for the same identity. This is heuristic: it flags conflicts in values, amounts, or dates that appear inconsistent across sources.

Using Contradiction Radar:

- Set From/To date range.
- Select an entity (ALL or a specific identity).
- Set Min confidence to ignore weak signals.
- Set Limit to cap the number of rendered items (performance control).
- Click Scan to compute the radar graph.
- Click items in the radar/table to view details; use Open Selected Evidence to validate in the original PDF page.
- Use Reset View if the radar is off-screen after navigation.

Do not treat a radar item as proof. It is an investigator prompt for where to look, not a conclusion.

3.9 Alerts and Patterns

Alerts and Patterns provides a triage workflow: rule-based alerts can be generated, reviewed, and exported.

- Use Search and Status filters (ALL/open/under_review/closed) to triage.



- Run Pattern Rules to scan recent passages and create new alerts (current rule set is intentionally small).
- Select rows and set status using Set Open / Set Closed.
- Export alerts as CSV or XLSX (written to Desktop/ariadne-bd/ariadne-output/).

3.10 Investigator Dossier

Investigator Dossier helps you package a case narrative and key findings into a portable bundle.

- Enter a dossier title/topic and click Create / Open.
- Curate notes and key items as you navigate other tabs (entities, evidence pages, alerts).
- Export as Markdown, HTML, or CSV for sharing and archiving.

3.11 Intelligence Report

Intelligence Report runs an iterative analysis loop over evidence passages and produces a structured narrative with recurring patterns and key points.

Inputs:

- Topic — what you want the report to focus on (recommended: one clear sentence).
- Language — auto / en / pt.
- Passages per step — controls how much evidence is processed per loop (effective range is bounded).
- Steps per loop — how many retrieval/extraction steps per loop.
- Total loops — how many loops to run (more loops = deeper synthesis, more cost if OpenAI mode).

Workflow:

- Fill the parameters and click Run.
- Monitor the log output while the report is being generated.
- When complete, review the report text.
- Use Export DOCX to generate a Word document in the output folder.

3.12 Settings

Settings is divided into General, AI, Operational Sources, Imagery Providers, and Storage. Use these pages to select AI options, configure provider credentials, choose source behavior, and review application paths.

- AI: select Local or OpenAI. AI Analyst uses the same AI selection saved here.



- OpenAI API key: enter the key when OpenAI is selected. After saving, the key is displayed masked.
- OpenAI model and reasoning effort: choose the model and reasoning level you want to use, then click Save. Use Test Selected AI before running a large analysis.
- Operational Sources: configure Aircraft, Vessels, Weather, and Infrastructure. Vessel mode supports AISStream.io only, Datalastic only, or AUTO.
- Imagery Providers and Storage: configure imagery-provider credentials, provider options, and storage locations used by supported imagery workflows.
- Datalastic is a commercial paid maritime-data provider. AUTO vessel mode can consume Datalastic credits, so use it only when that behavior is intended.
- Open-Meteo distinguishes free/non-commercial and Customer/commercial service modes. Choose the mode permitted for your deployment.
- Provider help buttons (?) explain how to obtain credentials and identify commercial/paid services where applicable.



3.13 Operational Globe overview

Operational Globe is an interactive geospatial workspace for operational intelligence. Enable the layers you need, define an Active AOI, refresh the relevant sources, and select objects on the globe to inspect them.

- Intelligence Layers: Satellites, Aircraft, Vessels, Weather, and Infrastructure are arranged in a compact two-column layer panel.
- Investigation Layers: Persons, Organizations, Locations, Events, and Relationships remain available separately from operational objects.
- Display: Labels is the global master switch for labels across satellites, aircraft, vessels, weather, infrastructure, AOIs, clusters, and other supported globe objects.
- Operational objects remain separate from investigative entities unless the user explicitly curates an observation or analysis into a Dossier.

3.14 Areas of Interest (AOIs) and geofencing

The Active AOI defines the main operational analysis area. Saved AOIs can remain visible simultaneously, but only one AOI is active for analysis at a time.

- CIRCLE: center latitude/longitude plus radius.
- RECTANGLE: two opposite corners.
- POLYGON: three or more vertices defining an arbitrary area.
- CORRIDOR: a centerline with a configured corridor width.
- Draw on Globe / Finish Drawing / Cancel Drawing: create advanced AOIs interactively on the Operational Globe.
- Import/Export GeoJSON and Import/Export KML: exchange AOI geometry with external GIS workflows.
- The Active AOI geometry is used to filter supported operational results. External providers can still impose their own coverage, availability, or geographic limits, which are shown in the relevant workspace when applicable.

3.15 Satellite intelligence and imagery

Satellite intelligence separates satellite activity, modeled AOI opportunities, confirmed imagery acquisitions, previews, and georeferenced products so that each type of information can be interpreted correctly.

- Enable Satellites in Intelligence Layers to display available satellite objects. Use the SATELLITE workspace to inspect selected satellites and analyze opportunities around the Active AOI.
- Mission information is shown when ARIADNE has a supported classification. UNKNOWN means that no reliable classification is available in the current data.



- AOI pass analysis shows modeled satellite ground-track opportunities near the Active AOI. These results do not confirm sensor coverage, tasking, collection, or image availability.
- Public imagery catalogues: Copernicus and USGS catalogue searches can return confirmed acquisition records, footprints, previews, and available public assets.
- Commercial imagery connectors: Planet, Maxar/Vantor, and ICEYE use provider-specific credentials and terms. Availability, ordering, download rights, and costs depend on the provider account.
- Before/After: search for acquisition pairs around an event/time, preview side-by-side/blend/difference views, and preserve the distinction between visual comparison and analytical change detection.
- Georeferenced products: when georaster support is available in your installation, ARIADNE can inspect compatible products, overlay them on the globe, and create registered comparison views.
- Use the available archive/download actions to preserve supported previews or local products where your provider rights allow it. Provider licensing controls what may be downloaded, retained, or redistributed.

3.16 Aircraft intelligence

Aircraft intelligence displays live ADS-B observations around the Active AOI and lets you review selected aircraft, recent tracks, AOI relationships, and timing around the selected Analysis UTC.

- LIVE: refresh aircraft for the Active AOI, optionally Observe Live, show labels, and select tracks.
- DETAILS: inspect ICAO24, callsign, registration when available, position, altitude, speed, track, recency, and AOI relation.
- CORRELATION: review local AOI entries/exits, closest observations, and timing around the selected Analysis UTC.
- Operational relevance ranking is transparent and based on observable factors. It is not a threat or intent score.
- Aircraft observations can be added to a Dossier without converting the aircraft into an investigative entity.

3.17 Vessel intelligence

Vessel intelligence supports AISStream.io and Datalastic. Choose the provider behavior in Settings -> Operational Sources -> Vessels before refreshing live vessel data.

- AISSTREAM.IO ONLY: uses only AISStream.io. A successful connection does not guarantee that vessel records are available for every area or at every time.



- **DATALASTIC ONLY:** contacts only Datalastic. Datalastic is a commercial paid API and successful queries can consume provider credits.
- **AUTO:** tries AISStream.io first and may use Datalastic only because the user explicitly selected AUTO. This mode can consume commercial credits.
- Provider choice persists after Save Settings and is reused on the next launch.
- **LIVE/DETAILS/CORRELATION** provide name, MMSI, IMO, callsign, vessel type, navigation status, SOG/COG, destination, local track history, AOI entries/exits, and operational relevance.
- Clicking a vessel on the globe displays the vessel name and AIS metadata in Overview when available; unnamed records fall back to callsign/MMSI.

3.18 Weather and environmental intelligence

Weather uses Open-Meteo weather and marine data around the Active AOI. The workspace is divided into LIVE, FORECAST, and ANALYSIS.

- **LIVE:** current/modelled temperature, humidity, precipitation, cloud cover, visibility, pressure, wind, gusts, and condition code/description.
- **Marine context:** wave height/direction/period, swell, sea-surface temperature, and ocean-current fields when available for the selected location.
- **FORECAST:** near-term hourly conditions for operational context.
- **ANALYSIS:** selects the model hour nearest Analysis UTC and derives an environmental-constraint score plus an optical-imagery suitability assessment.
- Open-Meteo free/non-commercial and Customer/commercial modes are distinct. Use the service mode permitted for your deployment.
- Weather/model data is context, not a certified local observation or safety forecast.

3.19 Infrastructure intelligence

Infrastructure uses OpenStreetMap/Overpass data to show public ground context around the Active AOI. Use Load / Refresh Infrastructure when you need current provider results for the selected area.

- **Categories:** Airports, Ports/Harbours, Energy, Transport, and ALL.
- Changing the category filters both the list and the globe immediately without issuing a new provider query.
- If the selected AOI is larger than the area supported by the current provider request, the workspace indicates the limitation. Narrow the AOI or reposition it if you need more focused coverage.
- **DETAILS/ANALYSIS** show source metadata, coordinates, distance/bearing, AOI relation, and nearest infrastructure by category.



- Infrastructure proximity is geographic context only and does not establish destination, affiliation, ownership, coordination, or intent.

3.20 Operational correlation

CORRELATION builds a unified UTC timeline from selected operational sources around Analysis UTC.

- Sources can include Aircraft, Vessels, Confirmed Imagery, Orbital Opportunities, Weather/Environment, and Infrastructure.
- Typical events include AOI entries/exits, nearest observations, confirmed acquisitions, environmental threshold entries, aircraft-near-airport events, and vessel-near-port events.
- Use Refresh Aircraft First, Refresh Vessels First, Refresh Weather First, or Refresh Infrastructure First when the relevant dataset is stale or absent.
- Open Related Object navigates back to the operational record where possible; Add Correlation Snapshot to Dossier preserves the selected analytical snapshot.
- Temporal proximity is not evidence of causation, coordination, tasking, ownership, or intent.

3.21 Patterns and anomalies

PATTERNS reviews available local operational history and is divided into ACTIVITY, ANOMALIES, and HISTORY.

- Choose a history window and object family, then click Analyze Local History. Use shorter windows for recent activity and longer windows when looking for repeated observations or longer-term deviations.
- ACTIVITY summarizes observations, AOI visits, entries/exits, sampled dwell, recency, recurring UTC hours, closest AOI distance, and infrastructure proximity.
- ANOMALIES can flag altitude, speed, movement, course, unusual observation hour, repeated AOI presence, or extended observed dwell deviations from local history.
- An anomaly indicates deviation from locally observed history only. It does not establish threat, intent, illicit activity, affiliation, ownership, or coordination.

3.22 Operational watchlists and alerts

WATCHLISTS lets you create persistent rules for observable Aircraft, Vessel, and Weather/Environment conditions. Rules are evaluated when relevant data is refreshed or when you click Evaluate Now.

- Aircraft/Vessel rules include object observed/reappears, enters Active AOI, and observed AOI dwell-time thresholds.



- Weather rules include wind-gust-above, visibility-below, precipitation-above, and wave-height-above thresholds.
- When a rule becomes true, a new alert appears in ALERTS. Acknowledge reviewed alerts and use Open Related Object when you need to return to the associated operational record.
- Alerts can be acknowledged, opened against the related object, added to a Dossier, or cleared after acknowledgement.
- After changing a rule or refreshing the underlying operational data, click Evaluate Now if you want to check the watchlist immediately.

3.23 Operational replay

REPLAY lets you review the historical operational observations available in ARIADNE for a selected UTC. It is divided into TIMELINE, OBJECTS, and EVENTS.

- Choose Replay UTC and a time window, then use previous/next, Play/Pause, speed controls, and the time slider.
- Aircraft and vessel positions are shown only where historical observations are available. If the selected replay time has a data gap, the missing object or interval remains unavailable.
- Historical trails can be limited to OFF, 15, 30, or 60 minutes and never reveal future track points relative to the replay clock.
- Replay can also display model weather context, mapped infrastructure, confirmed acquisitions, watchlist alerts, correlation events, and other available event records.
- Replay is a reconstruction from available/sampled data, not a certified continuous record of the real world.

3.24 Operational briefing

BRIEFING summarizes the operational data already loaded in ARIADNE without requiring an AI request. It can use ANALYSIS UTC or REPLAY UTC.

- BRIEF contains Situation, Timeline, Aircraft Activity, Maritime Activity, Satellite Activity, Environmental Context, Infrastructure Context, Alerts & Pattern Deviations, and Analytical Caveats.
- SOURCES identifies which inputs were available, absent, modeled, mapped, inferred, or locally observed.
- Export HTML and Export JSON create local briefing artifacts; Add Briefing to Dossier stores the structured snapshot.
- Review BRIEF and SOURCES before export. Missing or unavailable data remains identified rather than being presented as if it were present.



3.25 Encounter and proximity intelligence

ENCOUNTERS compares Aircraft and Vessel observations that occurred close enough in time to evaluate spatial proximity. It supports Aircraft <-> Aircraft, Vessel <-> Vessel, and Aircraft <-> Vessel.

- Configure horizontal separation thresholds and an observation time tolerance.
- Scope can require both objects inside the Active AOI, either object inside the Active AOI, or all local observations.
- Results include closest observed horizontal separation, paired UTC, observation time skew, observed encounter span, repeated-encounter count, and vertical difference for Aircraft <-> Aircraft where available.
- For long history windows, the interface may indicate reduced detail to keep the analysis responsive. Narrow the history window if you need a more focused encounter review.
- Encounter means spatial proximity between available observations. It does not establish communication, meeting, destination, coordination, affiliation, ownership, intent, or threat.

3.26 AI Analyst

AI ANALYST lets you ask natural-language operational questions and generate reports using the AI selected in Settings -> AI.

- ASK: enter a question about the Active AOI or current Replay context, select the source families you want to include, and click Analyze.
- REPORT: choose the report type, add any optional instructions, select the relevant source families, and click Generate Report.
- SOURCES: review the operational records referenced by the current AI analysis and open related objects where supported.
- Test Selected AI: confirm that the AI configuration saved in Settings is working before you run an analysis.
- Preview Context: review the operational data selected for the current request before sending it to the AI.
- Use the source-family checkboxes to control the scope of the analysis. Refresh or load the relevant operational workspace first if the data you need is missing.
- Click Analyze and wait while the progress indicator is active. The status area shows when the request is being processed or if an error needs attention.
- Read the response in the analysis panel, then use SOURCES to verify important factual statements against the underlying ARIADNE records.
- If OpenAI is selected, API usage is billable. For testing, start with a focused question, only the source families you need, and an appropriate reasoning level.



- If a response is incomplete, too broad, or not useful, narrow the question or selected data, check that the relevant operational sources are loaded, and run a new analysis.
- AI assessment is analytical output, not observed fact. Verify material claims against ARIADNE records and the original source where appropriate.

3.27 Operational source and provider summary

The following table summarizes the main operational providers available in ARIADNE. Provider availability, pricing, authentication, coverage, and terms are controlled by each provider and may change.

Capability	Default / supported provider	Credential / operational note
Satellite orbital data	CelesTrak	No separate provider key is normally required for standard public access. Use the source status/refresh controls when current data is unavailable.
Aircraft	adsb.lol / Airplanes.live	Refresh from the Aircraft workspace. The provider actually used is shown in the interface.
Vessels	AISStream.io / Datalastic	AISStream.io requires its own key. Datalastic is a commercial paid service. Choose AISSTREAM.IO ONLY, DATALASTIC ONLY, or AUTO in Settings.
Weather / marine	Open-Meteo	Choose the Free/Non-Commercial or Customer/Commercial mode permitted for your deployment.
Infrastructure	OpenStreetMap / Overpass	Use Load / Refresh Infrastructure. Public endpoint availability can vary.
Public imagery	Copernicus / USGS	Search supported public catalogues and use the available preview/open/download actions.
Commercial imagery	Planet / Maxar/Vantor / ICEYE	Provider-specific credentials, subscriptions, licences, and product rights apply.
AI Analyst	Selected AI from Settings	Uses the same AI selection saved in Settings. OpenAI usage is billed by OpenAI when OpenAI is selected.



4. Using the External Ingestion API

The API tab allows authorised external programs to send supported evidence directly to Ariadne. The API receives material; it does not collect data from external systems by itself. Collection, access authority, and transfer must be configured and controlled by the operator responsible for the external program.

4.1 What the API is for

- Receive camera captures, images, PDFs, e-mail, and WhatsApp material from authorised external programs.
- Associate incoming material with the integration that submitted it and display recent processing activity.
- Pass supported material into Ariadne's normal evidence-processing and indexing workflow.
- Make successfully processed material available to the same search, entity, relationship, timeline, location, face, AI, eDiscovery, and PREDICTIVE workspaces that use manually added evidence.

4.2 Before starting the API

- Confirm that Ariadne is running on the computer that will receive the material.
- Confirm that the external program is authorised to collect and send the intended information.
- Choose a clear integration name that identifies the program, installation, or data source.
- Decide which evidence types the integration is permitted to send: Camera / PDF, E-mail, WhatsApp, or an appropriate combination.
- Ensure that the receiving computer has sufficient storage and remains available while submissions are expected.

4.3 Starting and stopping the API

To make the API available:

- Open the API tab.
- Click Start API.
- Wait until API status shows STATUS: ON.
- Use Copy beside Endpoint when the authorised integrator needs the current connection address.
- Keep Ariadne open while external programs are submitting material.



To stop accepting new submissions, click Stop API and confirm that the status changes to STATUS: OFF. Stop the API whenever external ingestion is not required, particularly on shared or portable installations.

4.4 Creating an authorised integration

- Enter a descriptive Integration name.
- Select only the permissions required by that external program.
- Click Generate API Code.
- Click Copy and transfer the code securely to the authorised administrator or developer who is configuring the external program.
- Store the code securely. Ariadne displays the full code only when it is generated.
- Create a different code for every external program so that activity and revocation remain attributable.

The Authorised integrations table shows the integration name, a shortened key reference, permissions, status, and last use. The shortened reference identifies the credential without displaying the secret code.

4.5 Connecting an external program

- Start the API and copy the Endpoint.
- Generate or select the dedicated API code for the external program.
- Use Copy instructions to provide the current user-facing integration instructions to the authorised integrator.
- Provide the Endpoint, API code, selected permissions, and copied instructions through an approved secure channel.
- Ask the integrator to perform a controlled test using non-sensitive sample material before enabling routine submissions.
- Confirm that the test appears in Recent API activity and becomes searchable before accepting the integration for operational use.

4.6 Sending supported evidence

- Camera / PDF: provide the file and, where available, the capture date, source identity, location, event reference, and description.
- E-mail: provide a supported e-mail file or structured export containing the available sender, recipient, subject, date, message, and attachment information.
- WhatsApp: provide a supported export or structured message record containing the available participants, date, text, and media.
- Provide complete and accurate source metadata whenever possible. Missing optional metadata may not prevent receipt, but it can reduce provenance and analytical value.



- Avoid submitting the same source repeatedly. Use stable source names and identifiers so that later activity can be reconciled correctly.

4.7 Monitoring and confirming ingestion

- Review Client activity to see whether integrations are currently active or have submitted material recently.
- Review Recent API activity for the received time, integration, evidence type, processing status, and resulting file reference.
- Click Refresh if the Authorised integrations table does not yet show the latest use.
- Allow queued or actively processed material to finish before searching for it.
- Confirm important submissions through Smart Search, the relevant evidence workspace, or another standard Ariadne search before relying on them.
- If a received format cannot be fully indexed, preserve the activity record and review the original material through the available file reference.

4.8 Security, provenance, and revocation

- Treat every API code as a password. Do not place it in reports, screenshots, presentations, manuals, or ordinary e-mail.
- Grant the minimum permissions required for each integration.
- Right-click an integration row and choose Delete / revoke API key when access must end. Existing ingestion history is preserved.
- Revoke and replace a code immediately if it may have been exposed.
- Record the real source, capture time, responsible system, and collection authority outside Ariadne where required by the applicable evidence-handling procedure.
- Receipt by Ariadne does not establish authenticity, accuracy, lawfulness, consent, chain of custody, or evidential admissibility.

4.9 Common API problems

- STATUS: OFF: click Start API and wait for STATUS: ON before testing the external program.
- The external program is rejected: confirm that it is using the current dedicated code and that the required evidence permission is enabled.
- No recent activity appears: confirm that the Endpoint copied from the current session is being used and that the external program can reach the Ariadne computer.
- Material was received but is not searchable: wait for processing to finish, check the activity status, and confirm that the submitted format is supported for indexing.
- Unexpected duplicate activity appears: stop repeated submissions at the external program, retain stable source identifiers, and review the resulting evidence before continuing.





5. Using PREDICTIVE

PREDICTIVE investigates future-oriented questions against evidence already available to Ariadne. It searches selectively, follows relevant entities and documents, tests contradicting information, and returns a structured probabilistic assessment. It does not know the future and does not replace human judgement.

5.1 Suitable questions and required evidence

- Ask about a defined future event, recurrence, relationship, location, sequence, deadline, communication, organisation, document set, or other subject represented in the corpus.
- Use a measurable time horizon whenever the question permits one.
- Ensure that the database is current and contains enough dated historical evidence for the requested comparison.
- For legal or institutional forecasting, ensure that the relevant decisions, dates, court or body, subject matter, and comparable outcomes are actually present in the corpus.
- Use separate predictions for materially different events instead of combining unrelated questions in one request.

5.2 Selecting Subject / Context

- Leave Entire corpus / no fixed subject selected when the question should determine its own scope or concerns several entities.
- Select an existing person or entity when the prediction concerns a known subject.
- Type a subject directly when it is not listed but is clearly identifiable in the evidence.
- Click REFRESH ENTITIES after new evidence has been ingested or entity records have changed.

Selecting a subject narrows the initial context but does not prevent the research process from following relevant people, organisations, places, documents, dates, or relationships found during the investigation.

5.3 Research and forecast settings

- Deep Search OFF: use the AI selected in Settings for standard PREDICTIVE research.
- Deep Search ON: use the required OpenAI configuration for a more extensive iterative investigation.
- Max research iterations: set the maximum number of investigation cycles. This is a ceiling, not a target; research may stop earlier when sufficient evidence is available.
- Horizon: enter the number of days covered by the forecast.
- Use knowledge cut-off: enable this option when the assessment must use only information known by the selected date.



A knowledge cut-off is essential for historical questions such as: What was the probability on 15 May that this event would occur during the following 30 days? Later evidence must not be used to improve that earlier assessment.

5.4 Writing a predictive question

- State the event or outcome to be assessed.
- Identify the subject, entities, location, relationship, or sequence when these are known.
- State the time horizon or use the Horizon control.
- Add material conditions that must be distinguished from the main event.
- Ask for alternatives, contradicting evidence, or the information that would change the forecast when these are important to the decision.

Examples:

- What is the probability that Company X will miss another filing deadline within the next 30 days?
- Which known location is most likely to be used next by Entity Y during the next two weeks?
- What was the probability on 15 May that the negotiation would conclude within the following 30 days?

5.5 Running and stopping research

- Enter the question under ASK PREDICTIVE.
- Review Subject / Context, Deep Search, Max research iterations, Horizon, and any knowledge cut-off.
- Click RUN PREDICTION.
- Monitor Iteration, Searches executed, Documents reviewed, Relevant documents, Entities explored, and Current phase.
- Click STOP when the investigation should end. A model request already in progress may finish before the operation exits.

Research stops when the available evidence is sufficient, no materially new information is found, a technical limit is reached, the user stops it, or the selected iteration ceiling is reached.

5.6 Understanding the prediction

- Estimated probability: the central estimate for the defined event and horizon.
- Uncertainty: the reasonable range around the central estimate when the evidence supports an interval.



- Confidence: a qualitative assessment of evidential strength, consistency, coverage, and historical support.
- MODELLED PROBABILITY: calculated from sufficiently structured historical observations, comparable outcomes, frequencies, time patterns, graphs, or other measurable data.
- AI ESTIMATED PROBABILITY: an inferential estimate derived from documentary and relational evidence when a calibrated statistical model is not justified.
- INSUFFICIENT DATA, LOW CONFIDENCE, or INSUFFICIENT HISTORICAL SAMPLE: the available evidence does not support a stronger result.

A modelled probability and an AI estimated probability are not methodologically equivalent. Read the displayed type before comparing or communicating a percentage.

5.7 Reviewing WHY?, evidence, and the predictive graph

- Open WHY? to review supporting factors, contradicting factors, data gaps, and the stated method.
- Open DOCUMENTS to inspect the evidence selected for the final assessment. Double-click a row or use OPEN DOCUMENT.
- Verify decisive statements against the original document, date, page, metadata, and source context.
- Use PREDICTIVE GRAPH to review documented relationships, hypotheses, and predictions. Line patterns distinguish their evidential status; colour is not the only indicator.
- Treat preserved PREDICTIVE copies as references to their original source, not as additional independent evidence.

5.8 Backtesting

Backtesting asks Ariadne to make a forecast using only evidence available up to a historical cut-off, then compares the forecast with a later observable outcome when that outcome can be determined reliably.

- Use the same clear, measurable predictive question used for an ordinary forecast.
- Open BACKTEST and select the historical Cut-off.
- Click RUN BACKTEST.
- Review the forecast, actual outcome, outcome window, sample size, attempts, and coverage.
- Where shown, interpret Brier Score, hit rate, precision, recall, and false-positive rate together with sample size and coverage.



- Treat UNKNOWN / NOT SCORED as an unmeasurable or unresolved outcome, not as a successful or failed prediction.

Do not claim calibration or accuracy from a small or selectively chosen sample. Metrics are meaningful only when they derive from real, consistently defined outcomes.

5.9 What-if analysis

- Run or load the relevant prediction.
- Open WHAT-IF and describe one temporary change, such as excluding a document, lowering source reliability, or disregarding a disputed relationship.
- Click RECALCULATE SCENARIO.
- Compare the revised probability, alternatives, evidence, and confidence with the current result.

What-if analysis is temporary. It does not delete, alter, or reclassify the underlying evidence, entities, relationships, faces, events, or source files.

5.10 Anomaly detection

- Select or type a person or entity under Subject / Context.
- Open ANOMALIES and choose the number of baseline months.
- Click SCAN ANOMALIES.
- Review the anomaly type, date, observation, baseline, and evidence reference.

Anomalies may include unusual times or locations, new relationships, abrupt frequency changes, new sequences, co-occurrence changes, or the disappearance of a previously frequent pattern. An anomaly is a deviation from the observed baseline, not proof of intent, causation, wrongdoing, or future conduct.

5.11 History and evolving evidence

- Open HISTORY to review saved forecasts, probability type, confidence, status, and review state.
- Double-click a history row to load the available result.
- Compare probability history when new evidence has changed the assessment over time.
- Use a knowledge cut-off when the question concerns what Ariadne could have known at an earlier date.

A later conclusion does not erase an earlier conclusion. Review the evidence and confidence that were available at each recorded point in time.



5.12 Deep Search

- Configure a valid OpenAI API key in Settings and confirm that the required model access is available.
- Enable Deep Search only when the question justifies a more extensive investigation.
- Set a proportionate Max research iterations value and monitor usage, time, and progress.
- Allow the research process to follow relevant entities, relationships, dates, locations, snippets, and selected full documents.
- Review preserved relevant documents individually. Deep Search does not combine PREDICTIVE evidence into one evidence ZIP.
- Expect higher API use and longer completion time than standard PREDICTIVE research.

5.13 Responsible interpretation

- Use PREDICTIVE as a decision-support and research tool, not as an automatic decision-maker.
- For legal forecasting, distinguish legal merit from an observed institutional or jurisprudential pattern and verify every cited decision.
- Do not infer causation merely because entities or events repeatedly occur together.
- Do not convert documented warning signals into unsupported individual scores for crime, violence, self-harm, disease, or mental health.
- Record material limitations, missing evidence, disputed sources, and the precise question and horizon whenever a prediction is communicated.



6. Using eDiscovery

eDiscovery conducts structured external research, identifies and preserves relevant sources, and creates a reviewable dossier. It can also maintain a Living Dossier for later updates, claim testing, document additions, comparisons, and evidence products.

6.1 Purpose and scope

- Investigate people, companies, events, allegations, legal issues, factual questions, time periods, jurisdictions, and named sources.
- Search broadly while giving specified websites, databases, or official sources particular attention.
- Retain materially relevant sources and documents, organise findings, and produce a dossier with traceability information.
- Continue selected investigations through Living Dossier actions instead of starting an unrelated dossier for every development.

eDiscovery is a research and evidence-organisation aid. It does not replace disclosure obligations, legal hold procedures, privilege review, admissibility assessment, professional legal advice, or human source verification.

6.2 Required configuration

- Open Settings and configure a valid OpenAI API key.
- Select GPT-5.6 Sol and set reasoning effort to xHigh as indicated by the eDiscovery workspace.
- Return to eDiscovery and wait for OpenAI configuration verified.
- Confirm that OpenAI API key, Model, Reasoning effort, and Responses API each show a ready state.
- Ensure that network access, provider account access, billing, and any applicable spending controls are suitable for the planned research.

Create Dossier remains disabled until the configuration is verified and an investigation request has been entered.

6.3 Defining the investigation request

- Use What do you want to investigate? to describe the issue in plain language.
- Identify the principal people, organisations, events, documents, allegations, or questions.
- Define the relevant period, jurisdiction, and outcome required when these affect the research.
- Distinguish facts to establish, claims to test, and open questions.



- Use neutral wording where possible so that the investigation searches for contradictory and qualifying material as well as supporting material.

6.4 Specific sources and additional instructions

- Under Specific sources or websites to investigate, enter one website, domain, database, or page per line. This field is optional.
- Use Additional instructions to define date ranges, source priorities, jurisdictions, official-record requirements, languages, or output requirements.
- Specify primary or official sources when they should take priority over commentary or secondary reporting.
- Avoid imposing conclusions in the instructions. Ask the investigation to verify, distinguish, and cross-check material propositions.

6.5 Research parameters

- Maximum web searches sets the investigation ceiling. Higher limits may increase breadth, duration, and OpenAI API cost.
- Target relevant sources sets an approximate objective for materially useful sources; it is not a guaranteed count.
- Choose limits proportionate to the importance, breadth, and urgency of the request.
- Begin with a focused request and moderate limits. Increase them only when the first result identifies real gaps that justify further research.

Source quality, independence, authority, and relevance take priority over raw volume.

6.6 Creating and monitoring a dossier

- Review the request, specific sources, additional instructions, and research parameters.
- Click Create Dossier.
- Monitor Current status, the progress bar, Web searches performed, Sources identified, Relevant sources retained, Primary sources, Documents collected, and Sources discarded.
- Review the progress log for the current investigation phase.
- Click Cancel when research should stop. Completed recoverable work is preserved where possible.

Only one eDiscovery operation can run at a time. Wait for completion, cancellation, interruption, or an actionable error before starting another operation.

6.7 Reviewing a completed dossier

- Review the dossier name, completion date and time, search totals, relevant and primary source counts, document count, file location, archive size, and integrity reference.
- Click Open folder to inspect the dossier structure and individual files.



- Click Open ZIP to open the completed archive.
- Read the report and verify decisive claims against the retained primary sources and documents.
- Check dates, authorship, jurisdiction, version, publication status, and source independence before using a finding externally.
- Preserve the completed dossier and integrity information according to the applicable case and evidence-retention procedure.

6.8 Living Dossiers and recoverable investigations

- Click Refresh to update the lists of saved and recoverable investigations.
- Select a Dossier to review its version, claims, sources, documents, private documents, and open questions.
- Select a Recoverable investigation and click Resume investigation to continue preserved work after a supported interruption.
- Click Update Dossier to research developments, pending matters, and open questions. Add optional update instructions when a specific focus is required.
- Click Add Documents to add supported local documents to the selected Living Dossier.

Adding a private document keeps the local file as the canonical copy. If a later operation asks whether private documents may be used, read the prompt carefully. Selecting Yes may send eligible private content to the configured OpenAI service for focused retrieval; selecting No excludes that private content from the external AI-assisted operation.

6.9 Testing claims and generating work products

- Click Test a Claim, enter one clear assertion, choose a proportionate search limit, and review supporting, contradicting, and qualifying evidence.
- Click Generate Work Product and select the required product, such as a Legal Case Assessment, Statement of Facts, Draft Petition, Defence Assessment, Regulatory Complaint, Due Diligence Memorandum, Risk Memorandum, Board Briefing, Chronology, or Document Request List.
- Provide jurisdiction, parties, purpose, relief, audience, language, and other requirements when prompted.
- Treat every generated work product as a draft. Verify facts, authorities, quotations, procedural requirements, privilege, and professional obligations before use.
- Click Evidence Pack to create a local evidence package with the available hashes and traceability information.

6.10 Comparing dossiers

- Select two different dossiers under Compare.



- Click Compare Dossiers.
- Enter an optional comparison focus when a specific issue, period, party, or claim should receive priority.
- Review shared material, material differences, changes in claims or evidence, and any stated limitations.

6.11 Cost, privacy, and evidential practice

- OpenAI use may be billable. Set research limits deliberately and avoid repeating completed searches without a defined reason.
- Do not enter or upload material that the operator is not authorised to process or disclose.
- Apply privilege, confidentiality, personal-data, copyright, retention, and jurisdiction-specific review before sharing any dossier or work product.
- A retained web page or downloaded document is evidence of what was collected, not automatic proof that its contents are true.
- Record significant gaps, inaccessible sources, conflicts, and unresolved questions in the final review.



7. Data quality and schema management

7.1 Resolution: merging duplicates

Resolution is for cleaning identity duplication (e.g., “John Smith”, “J. Smith”, “Jonathan Smith”) so that the graph and analytics are less fragmented.

How to merge:

- Open the Resolution tab.
- Click Generate Candidates to compute likely duplicates.
- Select a candidate pair from the list.
- Click Preview Merge to see what will be combined (mentions, identifiers, evidence links).
- Choose which side is Primary (A is primary / B is primary).
- Click Merge to apply.

Rules: merges are allowed only for the same entity type. After merging, re-check Relationships and Timeline for improved consolidation.

7.2 Resolution: splitting conflated entities

Split is used when one entity incorrectly combines two different real-world identities (e.g., two people with the same name).

How to split:

- Select the entity to split (entity id) in the Split section.
- Click Load Split Preview to see the entity’s mentions and identifiers.
- Select which mentions and/or identifiers should move to the new entity.
- Provide a label for the new entity (be explicit: “John Smith (pilot)”, “John Smith (lawyer)”).
- Click Split to create the new entity and move the selected items.

After a split, review Relationships and Events Timeline: some relations may need manual interpretation until the corpus is re-indexed.

7.3 Ontology: active version, migrations, stats, validation

Ontology defines the formal schema used to interpret entities and relations (types, allowed relations, domain/range). Ariadne keeps the ontology versioned (semantic versioning) and can apply migrations.

Key actions:



- Reload — reload the active ontology from storage.
- Stats — compute dataset statistics by ontology class/relation.
- Validate — check whether stored relations violate domain/range constraints.
- Apply migration — upgrade the ontology to a newer version when available.

The default ontology includes PERSON, ORG, LOCATION, EVENT, DOCUMENT, ACCOUNT, and IDENTIFIER. Apply migrations only when they are supplied for your installed ARIADNE version, then review Stats and Validate after the migration.



8. Exporting and reporting

8.1 Alerts exports (CSV / XLSX)

In Alerts and Patterns, use Export CSV or Export XLSX to write the current filtered alert list to the output folder.

Notes:

- Exports include severity, status, rule name, description, and evidence reference (if available).
- Default output path: Desktop/ariadne-bd/ariadne-output/alerts_export.csv or alerts_export.xlsx.

8.2 Investigator Dossier exports (Markdown / HTML / CSV)

In Investigator Dossier, export curated content for sharing or archiving.

- Export Markdown — produces a portable, diff-friendly dossier (good for Git or audit logs).
- Export HTML — produces a viewable report for non-technical recipients.
- Export CSV — produces a structured extract (useful for downstream analysis).

All dossier exports are written under Desktop/ariadne-bd/ariadne-output/.

8.3 Intelligence Report export (DOCX)

In Intelligence Report, use Export DOCX to write the final report as a Word document.

Good practice:

- Include a clear topic statement (one sentence) and keep loops small when using OpenAI mode to manage cost.
- Treat the report as a synthesis draft: verify key claims by opening cited evidence elsewhere in Ariadne.

8.4 Operational Globe and AI exports

Operational workspaces provide additional export and Dossier actions for operational analysis, AOIs, briefings, imagery, alerts, replay, encounters, and AI-assisted outputs.

- AOIs: export GeoJSON or KML for circles, rectangles, polygons, and corridors.
- Operational Briefing: export standalone HTML or structured JSON.
- AI Analyst: export the current analysis as HTML or add the analysis to an Investigator Dossier for later review.



- Satellite imagery: archive supported previews/local products and manifests according to provider rights; georaster workflows can materialize local comparison products.
- Correlation, Weather, Patterns, Watchlists, Replay, Encounters, Aircraft, Vessels, and Infrastructure can add structured snapshots/observations to an Investigator Dossier.

8.5 PREDICTIVE and eDiscovery outputs

- PREDICTIVE lists the documents used for a result and allows supported evidence to be opened from DOCUMENTS.
- Deep Search may preserve relevant evidence as individual PDFs for later review. These files remain references to their original sources and must not be counted as independent duplicate evidence.
- eDiscovery creates a dossier folder and completed ZIP archive that can be opened from the result panel.
- Living Dossier actions may create updated dossier archives, claim-test results, work products, comparisons, or Evidence Packs. Use Open result when available.
- Review source rights, confidentiality, privilege, personal data, and integrity information before distributing any exported material.

8.6 Output folder

Default output folder is Desktop/ariadne-bd/ariadne-output/. It contains ariadne.log and supported exported reports/briefings. Additional operational data used by the application is stored under ariadne-storage/.



9. Keyboard and mouse reference

The exact interaction depends on the view type. The following is the current default mapping.

View	Pan	Zoom / reset
Geo Map	Middle mouse drag OR Alt + left drag. Buttons: Fit to Data, Reset View.	Mouse wheel OR +/- buttons.
Heat Map	Left mouse drag (hand pan).	Mouse wheel OR +/- buttons. Press 0 to reset zoom.
Events Timeline graph	Middle mouse drag OR Alt + left drag. Button: Reset Graph View.	Mouse wheel. Reset Graph View resets.
Contradiction Radar graph	Middle mouse drag OR Alt + left drag. Button: Reset View.	Mouse wheel. Reset View resets.
Evidence viewer (PDF)	Scroll to move page; use open controls in the viewer dialog.	Zoom depends on viewer implementation; prefer native PDF zoom controls if available.
Operational Globe	Left drag navigates the globe; click markers to select objects. AOI drawing uses successive left-clicks for geometry points.	Mouse wheel zooms. Use the workspace controls to manage layers. During AOI drawing, use Finish Drawing or Cancel Drawing to end the tool.



10. Troubleshooting

10.1 Evidence file not found

If Open Evidence reports that a file is missing, Ariadne could not resolve the original PDF path.

- Confirm the PDF still exists under files/ (same relative location as during ingestion).
- Avoid renaming/moving evidence after ingestion; if you did, reindex and keep paths stable afterwards.
- If the evidence is not a PDF, remember the evidence viewer renders PDFs only.

10.2 Geo Map shows “unresolved_or_ambiguous” frequently

- Increase Min score / Min confidence to reduce noise.
- Provide preferred countries via `ARIADNE_GEO_PREFERRED_COUNTRIES` (comma-separated) to bias ambiguous cities.
- Ensure `world_locations.json` is present and reasonably complete.
- Treat locations as hypotheses: open evidence and confirm context (e.g., “Braga” can appear in many non-location contexts).

10.3 OCR not running

- Install Tesseract OCR.
- Set `ARIADNE_TESSERACT_CMD` to the correct `tesseract.exe` path if autodetection fails.
- OCR is used only when the PDF page has no usable embedded text.

10.4 Ingestion is slow or appears stuck

- Large PDFs with OCR and face detection are expensive. Disable face search during ingestion if you only need text first.
- Prefer Reindex New/Changed Only for incremental updates.
- If you Cancel, keep “Resume interrupted build from checkpoint” enabled and re-run ingestion.

10.5 OpenAI / AI Analyst errors, access, or cost

- Use AI ANALYST -> Test Selected AI to confirm that the selected AI provider and model can respond before running a full analysis.
- If OpenAI is selected, verify the API key, selected model, account access, billing/quota status, and reasoning effort in Settings.
- OpenAI usage is billable. While testing, use a focused question, select only the operational source families you need, and choose a reasoning level appropriate to the task.



10.6 Backups and resets

Main data lives in Desktop/ariadne-bd/ariadne-storage/ariadne.db. Back it up before merges/migrations or before running large ingestion passes.

To hard reset the case state, close Ariadne and move the ariadne.db file aside (do not delete without a backup). Then re-open Ariadne and rebuild the database.

10.7 Operational Globe is blank, slow, or labels do not change

- If Operational Globe does not load correctly, restart ARIADNE and confirm that the application installation is complete. If the problem persists, restore or reinstall the application files.
- Disable unneeded Intelligence Layers if rendering many operational objects at once.
- Labels is the global display master. If a layer-specific label option is enabled but Labels is off, labels remain hidden.
- Infrastructure category selection filters both the list and map locally; use Load / Refresh Infrastructure only when the provider data itself needs refreshing.

10.8 Aircraft or vessel source returns no live records

- Check the Active AOI and provider status. A large AOI does not guarantee receiver/data coverage.
- Aircraft AUTO mode can try the available configured aircraft sources. Check the provider shown in the Aircraft workspace when reviewing live results.
- AISStream.io may connect successfully but still return no vessel records for a selected area. Use its Test Connection control and verify the Active AOI before concluding that the configuration is incorrect.
- For Datalastic, confirm the API key, subscription/credits, and any geographic limitation shown in the UI. The selected Vessel Provider Mode determines which service ARIADNE uses.

10.9 Weather or infrastructure source errors

- Weather: confirm the selected Open-Meteo mode is permitted for the deployment and that network access is available.
- Infrastructure: public Overpass endpoints are best-effort. Use Test Connection or configure another permitted endpoint if the selected endpoint is unavailable.
- After changing the Active AOI, refresh Weather or Infrastructure before using those datasets in Correlation or Replay if the workspace indicates that the existing data belongs to another area.



10.10 AI Analyst Analyze does not complete

- Run Test Selected AI first. If it fails, correct the AI settings or provider-account issue before using Analyze.
- Confirm that an Active AOI or Replay context is available, that at least one source family is selected, and that the operational data needed for the question has been loaded. Use Preview Context to check the selected data.
- If OpenAI is selected, remember that requests are billable. Review the selected model, reasoning effort, and provider-side spending controls before repeating an expensive analysis.
- Watch the progress indicator and status message. Wait for the response to complete or for an error message that identifies the action required.
- If the result is incomplete or unsuitable, narrow the question or source selection, check provider availability, and try again only when you are ready to submit a new request.

10.11 API does not start or no client activity appears

- Confirm that API status is STATUS: ON and that the Endpoint was copied from the current Ariadne installation.
- Confirm that the integration code remains active and includes permission for the submitted evidence type.
- Use a controlled test submission and review Recent API activity before enabling routine ingestion.
- If a code may be incorrect or exposed, revoke it and create a replacement for that integration.

10.12 PREDICTIVE is unavailable, incomplete, or slow

- For standard research, confirm that the selected AI in Settings is available.
- For Deep Search, configure a valid OpenAI API key and confirm supported model access in Settings.
- Use a focused question, realistic horizon, appropriate subject, and proportionate iteration ceiling.
- Confirm that relevant evidence has been ingested, indexed, and dated. PREDICTIVE cannot compensate for a missing historical sample.
- Use STOP if the operation is no longer required; allow any current model request to finish before expecting the state to return to READY.

10.13 eDiscovery Create Dossier is disabled or an investigation was interrupted

- Open Settings and confirm the OpenAI API key, required model, reasoning effort, provider access, and account availability.
- Return to eDiscovery and wait until the configuration is verified.



- Enter a non-empty investigation request. Create Dossier remains disabled without a verified configuration and request.
- If the operation is interrupted and appears under Recoverable investigation, select it and click Resume investigation.
- Reduce research limits or narrow the request only after reviewing the preserved progress and identifying the cause of the interruption.



Appendix A — Glossary

Entity: A node representing a person, organization, location, event, account, document, or identifier.

Mention: A text span in evidence that refers to an entity (e.g., a name occurrence in a PDF page).

Identifier: A structured token extracted from text (email, IBAN, phone, URL, domain, IPv4, NIF, hash).

Evidence page: A specific PDF page linked to a passage, mention, alert, timeline event, or face occurrence.

Relation: A link between entities (co_occurs / has_identifier / related_to).

Ontology: The formal schema describing allowed entity types and relations.

Migration: A versioned upgrade to the ontology schema.

Alert: A rule-generated item flagged for investigator review.

Operational Object: A satellite, aircraft, vessel, weather record, infrastructure feature, acquisition, alert, encounter, or similar operational record. Operational objects are not automatically investigative entities.

AOI (Area of Interest): A user-defined operational geometry used for filtering and analysis. Supported types include circle, rectangle, polygon, and corridor.

Confirmed Acquisition: An imagery catalogue record returned by a provider for a real acquisition/product. It is distinct from an inferred orbital opportunity.

Orbital Opportunity: A modeled satellite ground-track proximity event. It does not confirm sensor coverage, tasking, acquisition, or imagery availability.

Correlation: A unified UTC view that places records from multiple operational sources near the same analysis time. Correlation does not imply causation.

Pattern / Anomaly: A locally derived activity summary or deviation from ARIADNE-observed history. It is not a threat or intent classification.

Watchlist Rule: A persistent user-defined observable condition evaluated against refreshed operational data.



Replay: A reconstruction of available locally retained observations at a selected UTC. Gaps remain gaps unless explicitly identified otherwise.

Encounter: A proximity relationship between time-paired local observations of two operational objects.

AI Analyst: The workspace for asking operational questions or generating AI-assisted reports using the AI selected in Settings.

Source Record: An operational record shown in AI Analyst SOURCES so the user can review the data behind an AI-assisted statement.

External Ingestion API: The API workspace used to receive supported evidence from authorised external programs.

API Code: A sensitive credential issued to one authorised integration with defined evidence permissions.

PREDICTIVE: The workspace for evidence-based future-oriented questions, probabilities, backtests, what-if scenarios, anomalies, and prediction history.

MODELLED PROBABILITY: A probability calculated from sufficiently structured and measurable historical observations or comparable outcomes.

AI ESTIMATED PROBABILITY: An inferential estimate based on documentary and relational evidence when a calibrated statistical model is not justified.

Knowledge Cut-off: The latest date on which evidence may be treated as known for a prediction or historical assessment.

Backtest: A historical forecast made using evidence available up to a selected cut-off and scored only when a later outcome is measurable.

What-if: A temporary scenario that recalculates a prediction without changing the underlying evidence or Ariadne data.

eDiscovery Dossier: A structured research result containing retained sources, documents, findings, traceability information, and a completed archive.

Living Dossier: A persistent eDiscovery dossier that can be updated, supplemented, tested, compared, and used to generate later evidence products.

Evidence Pack: A local evidence package produced from a selected Living Dossier with the available hashes and traceability information.



Appendix B — Default ontology summary

Default ontology (v1.0.0) is intentionally minimal and is designed to be extended via migrations.

Class	Description
PERSON	Person (individual).
ORG	Organization / company / collective entity.
LOCATION	Location (city, country, address, etc.).
EVENT	Event (a node when meaningful).
DOCUMENT	Document (file) when represented as a node.
ACCOUNT	Account (bank, user, etc.).
IDENTIFIER	Identifier (base) – subtypes via IDENTIFIER * prefix.

Relation	Directed	Domain	Range
co_occurs	No	PERSON, ORG, LOCATION, EVENT, ACCOUNT, IDENTIFIER	PERSON, ORG, LOCATION, EVENT, ACCOUNT, IDENTIFIER
has_identifier	Yes	PERSON, ORG, ACCOUNT, LOCATION, EVENT	IDENTIFIER
related_to	No	PERSON, ORG, LOCATION, EVENT, ACCOUNT, IDENTIFIER	PERSON, ORG, LOCATION, EVENT, ACCOUNT, IDENTIFIER

Appendix C — Environment variables (advanced)

Most users will not need these options. They are intended for controlled deployments or data-location overrides.

Variable	Effect
ARIADNE_DESKTOP_DIR	Override the Desktop directory used for ariadne-bd/ and ariadne-output/.
ARIADNE_FACES_ROOT	Override the root directory for the face cache (faces.db, crops).
ARIADNE_GEO_CATALOG	Override path to world_locations.json



	used by Geo Map.
ARIADNE_GEO_PREFERRED_COUNTRIES	Comma-separated country codes/names to bias ambiguous location resolution.
ARIADNE_TESSERACT_CMD	Full path to tesseract.exe for OCR.

End of manual.