ChronoTracer and EDRM 2.0: Where It Fits

A ChronoTracer white paper by Richard Gorelick, Co-Founder and CEO. September 28, 2026.

Introduction

Sooner or later, every litigation and investigation turns on the same questions: who did what, with whom, and when. ChronoTracer is built for those questions. A technology platform delivered as a managed service, it organizes case evidence around discrete events and the people in them rather than documents, so legal teams can see what happened, from early case assessment through trial, with answers that trace back to the source evidence. Teams work with it through a web interface, and their AI platforms connect through an MCP server.

Because ChronoTracer works alongside the eDiscovery systems legal teams already have rather than replacing them, the question we hear most is a practical one: when in a case is ChronoTracer used? The Electronic Discovery Reference Model (EDRM) is the framework the industry uses to describe the stages of eDiscovery, and on September 1, 2026, EDRM released EDRM 2.0, the most significant revision in the model's twenty-year history. This paper answers that question using the new model: it explains what changed and why, and walks through where ChronoTracer fits phase by phase, while keeping the classic nine-stage vocabulary, still the language of most workflows and RFPs today, in view.

For readers who want the short answer: ChronoTracer plays a primary role in Processing (now part of EDRM 2.0's Data Acquisition group) and in Review, does its core work on the Analysis bar the new model makes continuous, and plays a supporting role in Presentation. It plays no primary role in the remaining phases. The phase-by-phase detail is in "Where ChronoTracer Fits," and the one-page version is the table in "Summary: ChronoTracer's EDRM 2.0 Footprint" at the end.

Background: The EDRM, and What 2.0 Changed

The Electronic Discovery Reference Model (EDRM) is the standard framework the legal industry uses to describe the stages of eDiscovery. George Socha and Tom Gelbmann created it in 2005 to answer two basic questions: what is eDiscovery at a practical level, and what are the basic steps you should consider taking? The first version was published in 2006, and it gave the industry what it was missing: a common language and a conceptual map for handling electronic evidence from beginning to end. It was never intended as a rigid, step-by-step procedure. It is a reference model: it identifies the key stages to think about and shows how they relate, and stages overlap, repeat, and feed back into each other throughout a case. Since its release it has been adopted by corporations, government agencies, law firms, and technology vendors.

The Nine Classic Stages

The classic model has nine stages, running roughly from left to right:

1. **Information Governance:** Policies and practices for managing ESI before litigation arises.
2. **Identification:** Locating potential sources of relevant ESI and determining scope.
3. **Preservation:** Protecting ESI against alteration or destruction (legal holds, etc.).
4. **Collection:** Gathering the ESI for use in the discovery process.
5. **Processing:** Reducing volume and converting ESI into forms suitable for review and analysis.
6. **Review:** Evaluating ESI for relevance and privilege.
7. **Analysis:** Evaluating ESI for content and context, including key patterns, topics, people, and discussion threads.
8. **Production:** Delivering ESI to other parties in appropriate forms.
9. **Presentation:** Displaying ESI before audiences at depositions, hearings, and trials.

A common way to read the diagram: as ESI moves from left to right through the model, its volume shrinks while its relevance rises.

Why the Model Was Due for Revision

Twenty years in, the EDRM remains the dominant framework for thinking about eDiscovery. It gives legal teams a defensible way to describe and manage the process, which is why almost every eDiscovery platform, service provider, and RFP in the industry still references its stages.

Practice evolved around the model, though. Structured data has always been part of the EDRM. Financial databases, phone logs, and similar sources have been in scope from the beginning, though in practice they typically follow a different path from the document workflow and are handled by a different team. What has strained the model in recent years is a different development: the explosion of conversational, distributed, and constantly changing data. Collaboration platforms like Slack and Teams, cloud storage, and mobile messaging produce evidence where the neat boundaries between "documents" break down, and where the most important facts often live in the connections between items rather than in any single item. And the next wave is already forming: as businesses put AI assistants and agents to work, the prompts, outputs, and actions of those systems are becoming discoverable evidence in their own right, an expected deluge that fits the document mold even less comfortably than chat data does.

Practitioners also pointed out that the model didn't account well for early case assessment (ECA) or risk analysis, both of which happen in practice before teams commit to full-scale processing and review.

What EDRM 2.0 Changed

EDRM 2.0 responds to those pressures. EDRM announced the project in March 2023, with a formal kickoff in May 2024; roughly 150 practitioners participated worldwide, led by co-project trustees Shannon Bales, Stephanie Clerkin, Rian Kennedy, and Brett Burney, with leadership support from David R. Cohen and Mary Mack. A draft circulated for public comment from June 30 to July 30, 2026, and the final model was released on September 1, 2026. EDRM describes the released model as integrating four elements:

Information Governance becomes foundational. Rather than sitting at the left edge of the diagram, the Information Governance Reference Model now anchors the EDRM, informing and governing every phase.

The left side consolidates into Data Acquisition. Identification, Preservation, Collection, and Processing are grouped into a single Data Acquisition framework, reflecting that in modern practice these steps often run concurrently and iteratively rather than as a strict sequence.

Analysis becomes continuous. Analysis is no longer a box in the line. It expands into continuous insight and feedback loops across the entire lifecycle, from identification through disposition.

Disposition joins the model. A new phase covers defensibly handling data after its use is complete: deletion, archiving, or return to normal retention.

The new diagram makes the changes visible. The Data Acquisition group sits at the left, feeding a flow that runs Review, Production, Presentation, and Disposition. Analysis runs as a bar beneath that entire flow, connected in both directions to the Data Acquisition group and to each phase downstream. And the Information Governance Reference Model spans the bottom as the foundation of the whole model.

Where ChronoTracer Fits

This section walks through the phases where ChronoTracer plays a role, then the phases where it does not. One thing holds throughout: ChronoTracer complements the platforms teams already run; nothing below replaces conventional processing, review, or production tooling.

Processing (within Data Acquisition): Strong Fit

Processing is where ChronoTracer departs furthest from conventional tools. The Processing phase is about reducing data volume and converting ESI into usable forms. ChronoTracer's deterministic pipeline is processing in that sense, but it performs a different conversion for a different purpose; it does not replicate or replace conventional eDiscovery processing.

A traditional eDiscovery processing platform takes raw data (email archives, file collections, etc.) and produces a reviewable document set: de-duplicated, indexed, and loaded into a review platform. ChronoTracer's pipeline takes that same raw data and transforms it into a structured, normalized database of discrete events. A single phone bill comes out the other side as thousands of individual call and text events, each tagged with date, participants, type, and source.

Most communication evidence is structure wrapped around content: an email is sender, recipients, and timestamps around a message body; a chat export is database rows whose payload is conversation. Conventional workflows carry all of it into review as documents. Much of the structure survives within each format, as threading, families, and chat rendering, but it is not normalized across formats into a single queryable layer of events. ChronoTracer works in the other direction, normalizing each item into discrete events (who, when, what kind) from the structure where it exists and from the content where it doesn't, as with a business letter or memo whose sender, recipients, and date the pipeline extracts. The downstream benefits of structure (querying, filtering, cross-referencing) then apply to evidence that would otherwise be read one document at a time.

Phone activity is a specialty: the pipeline has purpose-built parsers for carrier CDRs, phone bills, cell tower records, and mobile forensic artifacts, several different source formats that all yield the same call and text

events, plus identity resolution (associating phone numbers with known individuals). More broadly, the platform handles emails, texts (iMessage, SMS, WhatsApp, Signal), financial transactions, chat messages (Slack, Teams), calendar entries, and social media, with additional formats handled per engagement.

The pipeline is deterministic: same input, same output, every time. That matters here, where defensibility is the test. If opposing counsel challenges how evidence was processed, the repeatability of ChronoTracer's pipeline is one important part of the answer: the process can be re-run, examined, and validated, instead of being reconstructed from a language model's probabilistic output, as with tools that use LLMs to extract facts directly from raw files.

Review: Strong Fit, Different Approach

Traditional eDiscovery Review is, at its core, assessing documents for relevance and privilege. This is the most expensive and time-consuming phase of the EDRM, and it's where most of the industry's technology investment has gone (technology-assisted review in its generations: predictive coding, continuous active learning). EDRM 2.0 gives Review even more weight, describing it as the intersection where data volume, legal relevance, and strategic decision-making converge.

ChronoTracer doesn't replace document-by-document privilege review, and it doesn't position itself as a TAR tool. Instead, it changes the *unit* of review from the document to the event. Legal teams can search, filter, and explore a chronological event database instead of scrolling through thousands of individual documents. Review also becomes chronological in a way a document queue is not. Review platforms relate documents to other documents: threads, families, near-duplicates, conversation views. What they do not offer is event-level context across the full range of data types: the phone-log row, the transaction line, the calendar entry around a message. In ChronoTracer, a reviewer looking at a communication can pull up in one view what the same people were saying and doing in the same period, across the rest of the evidence in the database. And because the reviewer sets the scope, unrelated material stays out of the view.

ChronoTracer's AI tools, the built-in assistant and the MCP server, add natural-language querying and reasoning on top of the structured data. A reviewer can ask something like "show me the communications between Person A and Person B in the two weeks before the contract was signed" and get back actual events with source citations, with any AI narration grounded in those checkable records. The result is a more targeted kind of review, organized around the questions the legal team is actually trying to answer.

Analysis, the Continuous Thread: Core Strength

Analysis is about evaluating ESI for content and context: identifying key patterns, topics, people, and discussion threads. This is where ChronoTracer's value concentrates, and EDRM 2.0's treatment of Analysis is the most significant change in the model for ChronoTracer's purposes. In the classic model, ChronoTracer's analytical work had to be described as one stage plus spillover into the others. In 2.0, it is simply the Analysis bar: a continuous activity connected to the rest of the model, from data acquisition through disposition. That maps closely to how ChronoTracer works.

Because Processing has already normalized the evidence into an event database, Analysis queries run over structure, not over raw text. Users can see "who did what, with whom, and when" across all the evidence ChronoTracer has processed. The AI tools can break complex questions into multiple steps, and manual filters

(keyword search, Boolean queries, date ranges, participant filters) work alongside them, not in competition with them.

The chronological ordering itself is an analysis tool. Seeing events in sequence across different evidence types (an email thread, followed by a phone call, followed by a text message, followed by a financial transaction) reveals cross-evidence, event-level patterns that document-centric workflows do not surface efficiently, because they connect documents rather than the events inside them.

For investigations (criminal, regulatory, or internal), this is where ChronoTracer's value is most concrete: helping teams understand the narrative of what happened, built directly from the evidence. Because the structured event database persists across the life of the matter, the same analytical layer serves early case assessment, review support, formal analysis, and presentation preparation, wherever a team is in the lifecycle. That continuity is precisely what the Analysis bar in the new diagram depicts.

Presentation: Supporting Role

The Presentation phase involves displaying ESI before audiences at depositions, hearings, and trials. ChronoTracer isn't a trial presentation tool in the traditional sense, but the structured event views it produces are useful for trial preparation, deposition preparation, and building the chronological story of a case. The ability to quickly pull up every event in the database between chosen dates, involving chosen people, is the kind of retrieval litigators do when preparing witnesses or building opening statements. The final courtroom presentation happens through other tools; ChronoTracer builds the narrative those tools present.

Phases Where ChronoTracer Does NOT Play a Primary Role

The Foundation: Information Governance

In EDRM 2.0, information governance anchors the entire model, which means it also governs the platforms that hold evidence mid-lifecycle, ChronoTracer included. ChronoTracer is not an information governance tool: it doesn't manage data retention policies, legal hold workflows, or corporate records management, and it enters the picture after a litigation or investigation matter has already been identified. But it *operates inside the legal team's governance obligations*: generally single-tenant environments for per-case data isolation, secure transfer methods worked out per engagement, and agreed handling of the evidence when a matter ends. ChronoTracer has completed a SOC 2 Type II examination conducted by an independent third-party auditor, with no exceptions noted. And where governance requirements call for it, ChronoTracer now offers deployment on-premises or within the legal team's own cloud environment.

Inside Data Acquisition: Identification, Preservation, Collection

ChronoTracer doesn't help locate or scope potential sources of ESI. That work (identifying custodians, data sources, and the boundaries of what needs to be collected) happens upstream, typically through the legal team and their existing eDiscovery providers. Legal holds and data preservation are likewise upstream. ChronoTracer is a managed service that receives already-preserved, already-collected evidence through secure delivery methods worked out with each client; it does not perform the forensic collection itself.

Analysis can feed these calls, however. Working the event database can surface an unfamiliar phone number that keeps appearing around a key participant, a conversation other evidence refers to but the record doesn't contain, or a gap in a period that should be busy, each a signal that a source may exist that hasn't been identified or collected. The two-way arrow between the Analysis bar and Data Acquisition in the EDRM 2.0 diagram describes exactly this loop.

Production

ChronoTracer is not a production tool. It doesn't generate Bates-stamped production sets, load files, or handle the mechanics of producing documents to opposing counsel. Production happens through the existing eDiscovery platforms that ChronoTracer works alongside.

The decisions that drive a production are a different matter. Because ChronoTracer handles native files across formats and lets teams review that evidence in context, it can inform the responsiveness and scoping calls that determine what gets produced. In EDRM terms, though, that is Review and Analysis work; the production itself stays with existing platforms. EDRM 2.0 reflects this relationship with the two-way arrow between the Analysis bar and Production.

Disposition

Disposition, new in 2.0, covers defensibly handling data after its use is complete: deletion, archiving, or return to normal retention. ChronoTracer is not a disposition tool for the legal team's own data. It does, however, manage disposition of the evidence in its possession: when a matter ends, the handling of the evidence in ChronoTracer, whether return, deletion, or continued retention, is agreed with the legal team and carried out accordingly.

The Shape of the Footprint

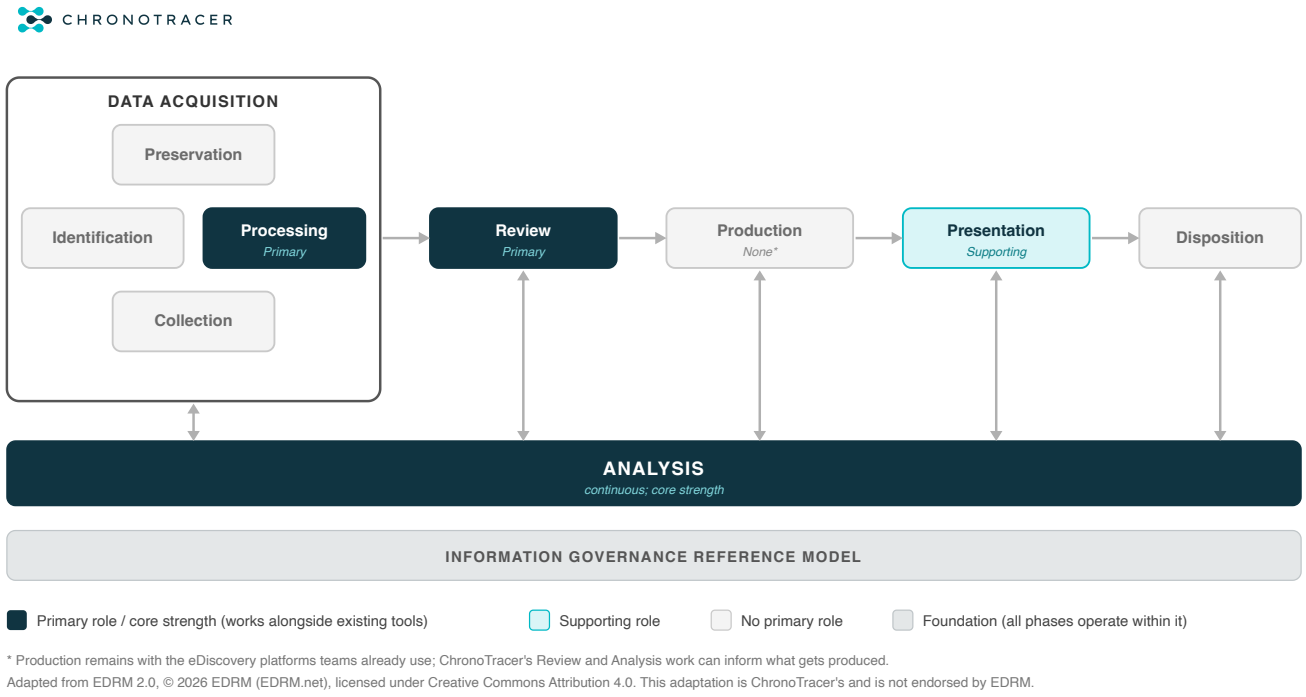

ChronoTracer's footprint, adapted from the EDRM 2.0 model, © 2026 EDRM (EDRM.net), used under CC BY 4.0.

The classic model made the answer easy to draw: ChronoTracer sat mostly on the right side of the diagram. EDRM 2.0 makes that answer more complicated. With the left side gathered into Data Acquisition and Analysis running as a bar beneath the whole flow, there is no simple right side anymore; the footprint has to be named phase by phase. Inside it: Processing (within Data Acquisition), Review, the Analysis bar, and Presentation. Outside it: Identification, Preservation, and Collection upstream, and Production and Disposition downstream. The split is deliberate. ChronoTracer adds an event layer alongside the platforms a team already runs, so bringing it into a matter does not mean replacing anything already in place.

AI in the Model: Grounding, Traceability, and MCP

Generative AI tools can answer confidently even when they are wrong or incomplete. For most business uses that is an annoyance; for evidence work it is disqualifying, because the entire point of discovery is that assertions must trace back to the record.

ChronoTracer's answer is built into the architecture, not promised in a policy statement. By design, the AI layer works from the structured event database the deterministic pipeline has already built rather than from raw files, and each answer traces back to the underlying source evidence, so it can be checked, not taken on trust. Structuring the evidence first closes off one class of silent failure that plagues AI over raw files: answers built from whatever the retrieval step happened to surface, incomplete in ways the reader cannot see. Teams reach that layer through the built-in assistant or through the MCP server, released in September 2026. MCP (the Model Context Protocol) is an open standard that lets AI platforms connect to outside tools and data

sources; ChronoTracer's MCP server is tested with Claude, ChatGPT, Gemini, and Microsoft Copilot, and designed to work with any AI platform that supports the standard.

The AI layer and the event database have distinct roles. The AI tools pointed at ChronoTracer can still make mistakes, as any AI tool can: a model can mischaracterize what it finds or miss something it should have found. What it cannot do is alter the evidence itself. Events enter the database only through the deterministic pipeline; the AI tools can annotate the record, applying tags, relevance levels, and comments, but cannot create, edit, or delete events. A citation therefore resolves by design to an event record, and the record links back to the source evidence it was extracted from; both the citation and the extraction can be checked.

In EDRM terms, this matters for two reasons. First, it is a concrete implementation of the idea EDRM 2.0 drew into the diagram itself: analysis as a continuous activity across the case rather than a single stage, the persistence described in the Analysis section above. Second, it offers one defensibility model for AI in discovery: deterministic structure underneath, AI on top, and a citation trail from every answer back to the evidence.

Summary: ChronoTracer's EDRM 2.0 Footprint

EDRM 2.0 Phase	ChronoTracer Role	Notes
Information Governance (foundation)	None	Operates within the legal team's governance obligations
Identification (Data Acquisition)	None	Outside scope
Preservation (Data Acquisition)	None	Outside scope
Collection (Data Acquisition)	None (receives collected data)	Secure transfer methods agreed per engagement
Processing (Data Acquisition)	**Primary (complementary)**	Deterministic pipeline builds the structured event database; does not replace conventional processing
Review	**Primary (complementary)**	Event-centric review with AI querying; works alongside existing review platforms
Analysis (continuous)	**Core strength**	Chronological, AI-assisted analysis across processed evidence types; serves the whole lifecycle
Production	None	Review and Analysis work in ChronoTracer can inform what gets produced
Presentation	**Supporting**	Builds chronological narratives for trial and deposition prep
Disposition	None (manages its own copy)	Handling of the evidence in ChronoTracer's possession is agreed per engagement

So, When in a Case Is ChronoTracer Used?

From early case assessment through preparation for trial. ChronoTracer structures evidence alongside conventional processing, changes the unit of review from the document to the event, does its core work in analysis, and builds the narratives other tools present at deposition and trial. Your review platform handles the documents. ChronoTracer handles what happened across them. In EDRM 2.0 terms, that means Processing within Data Acquisition, Review, and Presentation, with Analysis running across all of them. The approach scales to matters with tens of millions of documents. And each answer traces back to the source evidence.

In practice, ChronoTracer enters a matter as a managed service. The legal team delivers already-collected evidence through a secure transfer method agreed for the engagement; the pipeline processes it into an event database, generally in a single-tenant environment for per-case isolation; and the team works with the results through the web interface, while the AI platforms it already uses connect through the MCP server, alongside the review workflow already underway.

Acknowledgments

The author thanks Mary Mack of EDRM for feedback on an earlier draft of this paper. Any errors are the author's own, and the paper is not endorsed by EDRM.

Sources

- EDRM 2.0 Model (EDRM, September 2026; model © 2026 EDRM (EDRM.net), CC BY 4.0)
- EDRM 2.0 Project (edrm.net, September 2026)
- EDRM 2.0 Public Comment: Three Big Changes (EDRM, July 2026)
- EDRM Launches EDRM 2.0 Project (EDRM, March 2023)