


# From system records to study understanding

An Expert-led approach to protocol-aware, cross-system audit trail review

## Audit Trail Review as a Service

Clinical study records become more informative when their changes, timing and relationships can be examined together. ATRaaS is a managed service that combines preserved source evidence, study-specific interpretation and substantive Expert review, assisted by technology and AI.



### QA4Tech | White paper | September 2026

For clinical operations, data management, quality and sponsor oversight teams.

This paper describes QA4Tech's service and review method. Its scenarios are synthetic; the paper claims no customer results or measured performance advantage.

# Executive summary

A clinical study is a connected sequence of activities, while its electronic evidence is distributed across systems. Participant creation, questionnaire submissions, data corrections, approvals and study milestones may be recorded by different applications with different identifiers, structures and timestamp meanings. Reviewing each trail separately can leave relationships between those records unexplored.

Audit Trail Review as a Service (ATRaaS) is QA4Tech's Expert-led approach to that problem. It brings audit trail and supporting evidence into a shared study view, interpreted against the applicable protocol and critical data. Experts establish the review context, inspect records, formulate questions, assess signals, investigate alternatives and reach documented conclusions. Algorithms and targeted AI extend their analytical capacity.

The central proposition is a managed review service that connects study, technical, regulatory, system and assessment knowledge, with explicit scope and traceable reasoning. A combined dataset supports this work, but data integration alone does not establish that records mean the same thing, that an apparent discrepancy is real or that an export is complete.

## Three principles guide the approach

**Meaning before comparison.** Source records must retain their original definitions and evidence links. A shared format is useful only when it preserves the differences that matter to the review.

**Expert judgment throughout.** Review begins with study understanding and continues through direct examination, investigation and reporting. It includes questions initiated by Experts as well as signals generated by tools.

**Coverage as part of the conclusion.** The output should explain what was available, what was reviewed and what could not be established. Missing evidence and uncertain chronology must remain visible.

A bounded first engagement can test evidence readiness and review usefulness. It should include legitimate explanations and ambiguous evidence as well as known discrepancies. Progress should depend on the quality and usefulness of the work, rather than the number of alerts generated.

## In this paper

- The case for study-level review and its regulatory context.
- The expertise needed to interpret systems, assessments and study meaning.
- A method for connecting evidence and study meaning.
- Illustrative review scenarios and appropriate uses of AI.
- Independence, quality controls, the review report and a bounded first engagement.

# 1. Why the study is the relevant unit

The same participant may be represented in EDC, eCOA, IRT and other study systems. Each system can describe a different part of an event: the clinical activity, its data entry, a later correction, an approval or the propagation of a status to another application. A chronological list of technical records does not automatically explain the clinical-study sequence.

Consider a critical value corrected in EDC after a milestone recorded elsewhere. The review question might concern whether the correction matched an approved request, whether it affected data already used for a decision, or whether the observed sequence reflects delayed entry. Answering it can require the relevant source records, study expectations and an Expert's assessment of alternative explanations.

Cross-system review therefore changes the question from "What happened in this application?" to "What can the available evidence establish about this study activity?" Source-specific inspection remains necessary because it provides the detail on which the combined interpretation depends.

## Regulatory context, carefully bounded

EMA's 2023 guideline, section 6.2.2, describes documented, trial-specific, risk-based audit trail review focused on critical data, with proactive ongoing review expected unless justified. It contemplates both manual review and technology-assisted review of larger datasets. This supports a study-informed review approach; it does not establish a requirement to buy an external service. [1]

ICH E6(R3), sections 3.6.4-3.6.9, addresses documented transfers of activities and sponsor responsibility and oversight of service providers. Outsourcing review work therefore requires a clear division of responsibilities. [2] Applicable local requirements and study obligations must be assessed for the actual engagement.

## A defined addition to existing review

Each engagement begins by understanding the review already performed by the sponsor, CRO and other providers. Its potential contribution is a defined cross-system perspective, specialist investigation or additional review capacity. The extent of that contribution must be demonstrated on comparable evidence; the existence of a separate provider does not prove that current review is inadequate.

Reconstruction of the recorded participant journey can help explain a case. It should not silently expand the engagement into all clinical data review, monitoring or an independent GCP audit. Scope and outputs need to remain explicit.

## 2. Connected expertise in the review

A technical record and a study activity are different descriptions of an event. Understanding their relationship calls for several disciplines. The relevant question is whether the review can draw on those disciplines when interpreting evidence and investigating a case.

| Expertise the review draws on           | Study team | IT team | QA team | ATRaaS* |
|-----------------------------------------|------------|---------|---------|---------|
| Study & protocol knowledge              | ✓          | ●       | ●       | ✓       |
| Technology & audit-trail interpretation | ●          | ✓       | ●       | ✓       |
| Regulatory & data integrity knowledge   | ●          | ●       | ✓       | ✓       |
| Study-specific system knowledge         | ●          | ●       | ●       | ✓       |
| Assessment & instrument context         | ●          | ●       | ●       | ✓       |
| Interpretation across systems           | ●          | ●       | ●       | ✓       |
| Evidence-based investigation            | ●          | ●       | ✓       | ✓       |

**Team columns:** ✓ Typical core strength. ● Depends on skills and remit.

ATRaaS marks show the combined scope the service is designed to bring to each engagement. This is an illustrative comparison; actual team capabilities vary. Study, system and instrument knowledge is established during onboarding, with specialist input where needed.

### System knowledge and assessment knowledge are distinct

Knowing that a source is eCOA does not establish how its forms, permissions, timestamps or correction workflow are configured. Understanding the assessment also matters: which instrument and version were used, what the item measures, when it was administered and how the study uses it. A screening assessment and a later endpoint assessment can raise different review questions.

Experts establish this context during onboarding and investigate its implications in individual cases. Where a conclusion needs medical or other specialist judgment, define specialist input and escalation. ATRaaS does not assume the investigator's clinical decision-making or the sponsor's safety responsibilities.

### Coordinate the disciplines around the evidence

Study teams, IT and QA can bring these capabilities together internally. ATRaaS offers a managed way to coordinate specialist review and investigation across the agreed sources. The intended benefit must be evaluated on comparable evidence; neither a checkmark table nor outsourcing proves superior review quality.

## 3. A method built around review questions

The method begins with questions that matter to the study. Experts identify relevant critical data and processes, examine the applicable protocol versions and determine what evidence would be needed to assess those questions. The protocol is an interpretive source, not merely a document attached to an otherwise generic review.

### Establish an evidence and context baseline

The baseline should identify the systems and periods in scope, the exports received, their known limitations and the study documents used. Relevant Data Change Requests (DCRs), approvals, source definitions and system configurations may be necessary to interpret what the audit trails record. An export's existence does not establish that it includes every relevant event.

Experts should assess the effect of missing information early. Some questions may remain answerable with a stated limitation; others may need additional evidence or a revised scope. These decisions belong in the review record so that later conclusions can be understood against the same baseline.

### Preserve meaning during transformation

A common event representation makes filtering and comparison practical. It should preserve original fields, values, codes and source references, while distinguishing any derived interpretation. A questionnaire submission containing many answers should not be presented as many independently timed observations if the source provides only one submission event.

Study/system transformation rules should be recorded and versioned. Where a code has a historical meaning, that definition must accompany its interpretation. A mapping can be reused, but protocol amendments and source changes can require reassessment. "Mapped once" is not evidence that a source can be processed indefinitely without review.


### Reconcile identity and time explicitly

Linking records across systems requires a documented basis for the participant, visit, form or item relationship. Ambiguous matches should remain visible. Similarly, activity time, data-entry time, submission, server receipt and audit-record creation are distinct concepts. Sorting them together without qualification can create a misleading sequence.

The combined view is therefore an evidence-supported interpretation. It should retain access to the original records and the uncertainty associated with each relationship. Experts can challenge a mapping or request further analysis when it does not explain the source evidence adequately.

## 4. Evidence flow and the Expert workspace

The review flow separates source evidence, study context, analytical assistance and Expert conclusions. Experts contribute before, during and after automated analysis.



A participant timeline can support this process by showing the recorded journey across sources. Experts should be able to inspect event details, later changes, applicable protocol context and unresolved links. The journey can be shown in source swim lanes or as a combined chronology. Neither layout changes what the evidence can prove.

The workspace also supports direct review of unflagged activity and Expert-initiated questions. Investigation is iterative: an observation may lead to a targeted query, further source information and a revised interpretation before a conclusion is documented.

## 5. Three illustrative review questions

The scenarios below are synthetic. They demonstrate reasoning, not observed findings or claims about particular vendors.

### A critical correction and its change request

An audit trail records a post-submission change to an efficacy questionnaire item. The Expert establishes the instrument and version, item meaning and use at the relevant study stage. A DCR describes a correction to an apparently related record. The first task is to establish whether the participant, visit, form version and data item actually match. The Expert then compares the requested and recorded changes, including scope and timing where available.

Possible explanations include a correctly implemented correction, an incomplete request, a change to the wrong record, a later amendment or missing evidence. AI may assist the comparison, but the existence of an approved DCR does not by itself prove that the recorded change was appropriate. A useful conclusion identifies the supported relationship and any unresolved discrepancy.

### A questionnaire before an expected milestone

An eCOA event appears before an IRT randomization record. A simple ordering check can identify the apparent difference, but an Expert must establish whether the questionnaire was a permitted screening activity or one that the protocol required after randomization.

The timestamps also need interpretation. A recorded submission time may differ from completion time; a server update may occur later than the underlying milestone. If the identity or chronology cannot be established, the outcome should be an unresolved ordering question rather than a definite protocol finding.

### A change near a study decision point

Several critical values change shortly before an analysis cutoff. The concentration of activity may merit examination, but timing alone does not establish misconduct. It could reflect a planned data-cleaning exercise, a migration, a legitimate batch correction or an issue requiring investigation.

The Expert examines the relevant actors, reasons, approvals, values and study context. A cross-system perspective may help determine whether the same underlying activity is reflected consistently elsewhere. If the review cannot obtain essential source evidence, the limitation belongs in the conclusion.

### What these scenarios demonstrate

Signals identify questions to pursue. Findings require supporting evidence and a documented assessment. A legitimate explanation and an inconclusive outcome are both necessary parts of a useful review method; they prevent apparent anomalies from becoming unsupported claims.

## 6. How technology and AI assist Experts

Different analytical tasks require different methods. Explicit timing rules, repeated-event checks and defined comparisons can often be implemented deterministically. Contextual questions may benefit from AI-assisted extraction or comparison. Neither approach removes the need to establish the meaning and completeness of the evidence being analysed.

### Deterministic analysis

Rules can help identify known relationships, such as a change after a specified milestone or a mismatch between mapped values. Their usefulness depends on correctly defined prerequisites. If the milestone is missing from the export, the rule should indicate that the question could not be evaluated for that scope; it should not silently record a pass.

Rules should be tied to the relevant study expectations and data versions. Experts need to understand why a signal was raised and what records it used. Changes to mappings or rules may change the result, so the version used in a completed review should remain identifiable.

### Targeted AI assistance

AI can help organize a bounded set of evidence, compare a DCR with recorded changes, identify questions in a reason-for-change statement, or prepare a draft summary with source references. Each use should be evaluated against representative cases and known limitations before it is relied on in a review. AI is used case by case, and permitted use is agreed for each engagement.

The relevant evidence should be selected and preserved so that an Expert can verify the output.

Unsupported explanations, omitted context and inconsistent interpretations remain possible. The use of AI should not obscure the underlying records or make a suggested relationship appear certain.

### Substantive Expert review

Experts establish priorities, examine evidence directly, initiate investigations, request further analysis, weigh alternative explanations and decide what belongs in a report. They can challenge automated suggestions or pursue a concern that produced no signal. Review of unflagged activity is a method-design question, not something to eliminate by default.

The result of this division of work should be assessed through detection quality, missed cases, false signals and traceability. Faster processing is not sufficient evidence of better review, and agreement between two AI systems is not independent confirmation of correctness.

## 7. Independence, quality and evidence limits

Independent review can provide a useful separation between the people assessing evidence and the parties generating or managing it. In ATRaaS, independence comes from QA4Tech's position: it is independent of the sponsor, CROs, sites and the vendors whose systems produce the records. That independence is supported by arrangements for conflicts, access to evidence, disputed interpretations and reporting that are agreed for each engagement. It is not a claim that all ATR must be externally performed.

### Independence needs an operating model

A sponsor may engage the service directly, or a CRO may subcontract specialist review. Either arrangement requires clarity about who receives conclusions, how disagreement is recorded and when a matter is escalated. A contractual relationship should not determine the substance of an Expert's assessment.

Independence also does not establish upstream integrity. A downstream reviewer may receive an incomplete or altered export. The review needs an account of provenance, extraction scope and available reconciliation evidence, with explicit limitations where independent confirmation is unavailable.

### Proportionate assurance before reliance

ICH E6(R3), section 4.3.4, describes validation based on intended use and risk, with appropriate validation before use and consideration of subsequent changes. [2] A small engagement or an Expert's final decision is not, by itself, evidence that its supporting tools are suitable.

For ATRaaS, readiness should address the transformations, identity and time interpretation, review methods and reporting functions on which conclusions depend. Testing should include missing inputs, ambiguous evidence and legitimate exceptions as well as expected detections.

### Data handling and controlled change

A confidential disclosure agreement is in place before any evidence is received; the export format, a secure delivery route such as SFTP and the delivery frequency are then agreed. Access, permitted AI use, retention, recovery and exit arrangements should also be defined for the engagement. Evidence containing unblinded information needs particular consideration when assigning review access. These arrangements are agreed in the engagement documents.

An amendment, revised export or corrected mapping may affect previous interpretations. The operating method should record what changed, assess whether earlier work needs revisiting and preserve the basis of previously delivered conclusions. Repeatability requires this history, not just the ability to run the same script again.

## 8. Starting with a bounded first engagement

A first engagement should test whether the service can answer useful questions with traceable reasoning. It also reveals evidence and method limitations before broader use.

### Begin with evidence readiness

Select one study or representative dataset and a few relevant sources. Define the questions, outputs and expertise needed, including system and instrument knowledge and any specialist input. Assess the exports, context, identity links and timestamps, then identify missing information. Set a decision point to proceed, revise scope or stop with a documented readiness output.

### End with a review report and a decision

A pilot review covers one study, a small number of its systems and one or two agreed questions. It ends with its review report and a discussion of conclusions, unresolved matters and follow-up, after which the customer decides whether to stop, refine or extend.

A retrospective review can show feasibility without establishing prospective performance. Study mix, changing exports, reviewer experience and investigation demands may affect later results. No performance targets are established in this paper.

## 9. From review work to useful decisions

The service delivers an understandable account of the evidence and the Expert's assessment. A report should distinguish observations, supported findings, legitimate explanations and unresolved cases. It should identify the sources and periods reviewed, the relevant study context, the methods applied and the limits of the conclusion.

Evidence-linked case records can support customer follow-up by making the reasoning inspectable. Recommended actions need an agreed owner and escalation route. The service does not claim to resolve independently every clinical or operational matter identified during review.

### The central value hypothesis

Protocol-aware review across systems may reveal relationships that are difficult to assess in separate trails. Experts can use that perspective to formulate better questions, investigate potential inconsistencies and develop more useful explanations. The degree of improvement must be demonstrated through representative review work.

ATRaas combines that hypothesis with a managed delivery model: specialist review, supporting technology and documented outputs within a defined engagement. Every engagement begins with an evidence readiness assessment or a focused pilot – a targeted investigation begins with a readiness assessment scoped to its concern – and expands only where the evidence supports doing so.

**The first question is practical: what does your study team need to understand, and what evidence would allow an Expert to reach a defensible conclusion?**

**Discuss a bounded first engagement:** [info@qa4tech.com](mailto:info@qa4tech.com) | [qa4tech.com/audit-trail-review](https://qa4tech.com/audit-trail-review)

### References and evidence status

**[1] European Medicines Agency.** Guideline on computerised systems and electronic data in clinical trials, 9 March 2023, EMA/INS/GCP/112288/2023, section 6.2.2, p. 20. [Official guideline.](#)

**[2] International Council for Harmonisation.** ICH E6(R3) Guideline for Good Clinical Practice, final Step 4 version, 6 January 2025; sections 3.6.4-3.6.9, printed p. 25, and 4.3.4, printed pp. 49-50. [Official guideline.](#)

Sources checked 24 September 2026. These guidelines provide context within their applicable scope; they do not endorse ATRaaS. The method and examples are QA4Tech's own. This is a conceptual white paper, not an empirical performance study or a complete regulatory assessment.