

MULTI-SYSTEM CASE INTELLIGENCE

AI CaseQuery Connect

Firms often keep current and legacy case-management systems running side by side for months or years. CaseQuery Connect gives the firm one reporting and query layer across them - without consolidating databases, changing systems of record, or waiting until older platforms can be retired.

One question across Case Systems 1, 2 & 3

One reporting layer across current & legacy systems

Duplicates resolved before totals

Read-only connections · local execution

WHEN CASE SYSTEMS HAVE TO COEXIST

The firm changed systems. The work did not.

As firms transition between case-management platforms, older systems often remain in use for months or years before they can be fully retired. Active, historical and recently converted matters may live across multiple databases. Operations still needs one answer for client reporting, workloads, deadlines and management questions - even while the underlying systems disagree about names, status codes and where a matter lives.

*"How many active matters does the firm really have - without counting the same case twice?"**"Show every file with a sale date in the next 30 days, regardless of which case system holds it."**"Which matters still live only in a legacy system?"**"Where do the current and legacy systems disagree on status, attorney or important dates?"*

One Business Vocabulary

CaseQuery maps each system to common meanings such as Case, Client, Status, Referral Date, Sale Date and Attorney — so users do not need to know three schemas.

Safe Multi-System Totals

If overlap is possible, CaseQuery resolves matter identity before a count, sum or distinct total is presented. A copied matter is not silently counted twice.

Systems Stay Independent

Current and legacy platforms remain separate systems of record. CaseQuery queries each where it sits and combines only what can be safely combined.

Source Always Visible

Answers retain the source breakdown and any coverage gap. If one system is unavailable or a concept is unmapped, CaseQuery does not disguise a partial answer as a firm-wide total.

HOW IT WORKS

One question. Three systems. One controlled query plan.

The model understands the business question once. CaseQuery's own code - not the model - decides how to execute it safely against each connected system.

1 Connect each system read-only

Case System 1, Case System 2 and Case System 3 remain independent. CaseQuery runs ordinary read-only queries against each; no cross-database linked server is required.

2 Translate to one canonical meaning

Different table names, status codes and business fields are mapped to the firm's shared vocabulary and business rules during setup.

3 Choose the safe execution path

If the systems hold different matters, results can be combined directly. If the same matter may appear twice, CaseQuery requests identity keys first, resolves overlap, then aggregates.

4 Return the answer with coverage

The firm sees the consolidated answer, the contribution from each source, and any limitation that would make a firm-wide total unsafe.

Question	"How many active foreclosure matters does the firm have?"
Canonical meaning	Active foreclosure + connected source scope + firm definition of active
Planner decision	If the same matter may exist in more than one source, retrieve identity keys first; otherwise aggregate safely at the source
Sources	Case System 1 · Case System 2 · Case System 3
Answer	One firm-wide result, with source and coverage details retained.

Built for systems that have to coexist

Replacing a case-management platform does not make the older database disappear. CaseQuery Connect is built for the period when current and legacy systems must remain in service together.

- ✓ Query current and legacy systems together while both remain in service
- ✓ Keep each database independent - no linked servers or cross-database views required
- ✓ Continue firm-wide reporting before the legacy platform can be sunset
- ✓ Add or retire a source without redesigning the reporting layer

Business meaning stays constant. "Active," "sale scheduled" or "first legal" can have one firm definition even when each source implements it differently.

Identity is resolved before anything is totalled

The dangerous design asks all three systems for *COUNT(*)* and adds the answers. Once duplicates are buried inside those numbers, nothing downstream can find them.

Deterministic matches can be automatic

Prior-system reference, prior file number, loan number plus client, servicer file number or another trusted cross-reference can establish that two rows represent the same matter.

Probable matches are never silently merged

Borrower, property, dates or similar clues can create a review candidate, but a possible match remains separate until a person confirms it. Refilings and second liens must not disappear into a fuzzy merge.

THE CONNECT LAYER, NOT A NEW CASE SYSTEM

Leave the case data where it already lives

CaseQuery Connect does not become another system of record. The source databases remain authoritative. Connect maintains the mappings, business definitions and narrow identity index needed to recognize the same matter across sources, then runs the actual queries locally against the systems that own the data.

3 case systems can participate in one business question

1 canonical vocabulary for the AI and the firm's rules

0 need to expose three raw schemas to the user asking the question

Primer shows coverage before users rely on it

A connected database is not automatically equivalent to the others. An older system may contain cases and status but not the same fee, document or assignment detail as the current system.

Business concept	Case 1	Case 2	Case 3
Cases / status	✓	✓	✓
Client	✓	✓	✓
Sale date	✓	✓	✓
Assigned attorney	Map	✓	✓
Fee detail	—	✓	✓

A missing concept is shown as a coverage limitation - not quietly omitted from a "firm-wide" answer.

Cross-system reconciliation shows how the systems overlap

Before unified totals are enabled, a reconciliation pass can establish which matters appear in more than one system and where mappings need calibration.

- ✓ Which matters appear in each source
- ✓ Deterministic matches and unmatched matters
- ✓ Status, client, attorney and important-date disagreements
- ✓ Client-name variants and status-code mapping issues
- ✓ A reviewed identity map for safe firm-wide reporting

A partial answer is not presented as a total

If Case System 2 is offline, an identity relationship has not been established, or one source does not contain the concept being asked about, CaseQuery can return the available source breakdown and explain the gap - but it does not manufacture a single firm-wide number.

Case System 1 Case System 2 Case System 3

- Deduplicate before COUNT or COUNT DISTINCT
- Rewrite averages using the underlying sum and count
- Apply source precedence only after identity is established
- Keep per-source provenance through the final answer
- Refuse a consolidated result when it cannot be computed safely

HOW CASEQUERY CONNECT GETS TO ONE SAFE ANSWER

