Docket

DOCKET / RESEARCH & PRACTICAL GUIDANCE

# Acquisition Target Research Update Frequency in 2026

> Build a more reviewable acquisition pipeline with Docket. Our deal-origination software and managed research help private equity teams screen companies against a mandate and understand the evidence behind each finding.

Published by Docket · 2026-09-22 · acquisition target research update frequency · acquisition target data freshness · private equity research · target monitoring · due diligence · research operations · data quality · m&a target research · docket

Original article: https://docket.capital/articles/acquisition-target-research-freshness



### In this article

1. Executive Summary
2. Introduction and Background
3. Key Changes
4. Implementation Considerations and Process Changes
5. Freshness Policy Matrix
6. Data Analysis and Evidence
7. Case Studies and Real-World Examples
8. Implications and Future Directions
9. Frequently Asked Questions (FAQs)
10. Conclusion

## Executive Summary

Acquisition target research should not be governed by one **last updated** badge. The defensible unit is the field, and each field needs four clocks: when the fact became effective, when the source published it, when the research system collected it, and when an analyst reviewed it. Those clocks can diverge materially. The U.S. Securities and Exchange Commission (SEC) says its submissions data are updated throughout the day in real time (sec.gov), yet its bulk archive is republished nightly at about **3:00 a.m. Eastern Time** (sec.gov). Companies House can push changes as they occur (developer-specs.company-information.service.gov.uk), while its month-end snapshot can arrive within **five working days** (gov.uk). “Real time” therefore describes an access path after a reported event. It does not prove that the underlying company fact is current.

The recommended acquisition target re-research service level agreement, or **SLA**, combines volatility with decision consequence. Legal status, ownership and the company domain are high-consequence identity fields. They should receive event monitoring where an authoritative stream exists, plus a short maximum research age before screening, handoff or outreach. Headcount, services, geography, revenue estimates and contacts need different clocks because their sources have different meanings and lags. The Bureau of Labor Statistics (BLS), for example, publishes Quarterly Census of Employment and Wages data within **six months** after a quarter (bls.gov), and it does not publish an individual company’s employee count (bls.gov). That makes QCEW useful context, not a current target-level headcount source.

Operationally, a target record should carry **current, stale, unknown, contradicted, or revalidated** at field level. Stale means the policy age has expired, not that a fact is false. Unknown must remain distinct from zero: European Union data-quality guidance defines a null as an unknown value (op.europa.eu). Contradictions should remain visible until resolved. Accuracy is also separate from timeliness because accuracy asks whether data matches reality (gov.uk).

The practical answer to “how often” is therefore: refresh on relevant events first, schedule reviews only to cover gaps, and force re-research at consequential decision gates. Track **stale high-risk fields divided by all high-risk fields**, together with median and 90th-percentile age by field. Estimate workload as records

multiplied by fields due multiplied by a user-entered minutes-per-field assumption. No cadence guarantees correctness. The policy controls observable age and review discipline while retained evidence, explicit unknowns and contradiction handling control what a team can verify.

## Introduction and Background

An acquisition target record is not a document with one natural expiry date. It is a bundle of claims about identity, ownership, operations, location, scale and reachability. A legal status drawn from a registry, a service line taken from a website, a revenue estimate derived from public clues and an email address confirmed by delivery feedback can all be collected on the same day without being equally current. **Acquisition target data freshness** is therefore the fitness of each claim's age and provenance for a particular decision.

That definition follows a broader data-quality principle. The World Wide Web Consortium (W3C) defines currentness in relation to the "right age" for a specific use (w3.org). UK government guidance likewise defines timeliness by whether the lag from collection to availability suits the intended use (gov.uk). For private-equity origination, "intended use" changes from broad market mapping to target screening, partner review, handoff and approved outreach. A fact reusable for landscape analysis may be too old for a named-target approach.

Docket is a direct provider of sourced acquisition target research through its Triage, Scout and Audit agents. Its first-party description says each researched answer includes a source, supporting excerpt and collection date. Its research library also says its guides link to sources and explain evidence limits. Those practices are necessary for freshness operations, but not sufficient: a collection date does not reveal the source's effective date, publication lag or subsequent change.

This report addresses the practical buyer question: **when should company research be reused, refreshed automatically, or re-researched?** The answer is a field-and-decision policy. It separates observable timestamps, ranks fields by consequence and volatility, uses events where sources expose them, applies scheduled reviews where event coverage is incomplete, and keeps uncertainty visible.

## Key Changes

### Replace one timestamp with four clocks

The first change is semantic. A record should never overload `updated_at` with several different events. Four evidentiary timestamps answer four different questions; `next_review_at` is a policy deadline rather than an evidentiary clock:

- `as_of`: When was the fact true or effective according to the source?
- `source_published_at`: When did the source expose or formally issue the record?
- `collected_at`: When did the research system retrieve and retain the evidence?
- `reviewed_at`: When did an analyst decide that the evidence supported the field value?
- `next_review_at`: When will policy require another check if no triggering event arrives?

The distinction is not theoretical. Delaware describes an effective date separately from the day its office received a filing (icis.corp.delaware.gov). Dublin Core similarly distinguishes formal publication from the date a resource changed (dublincore.org). Collection and review are local workflow facts, not substitutes for either source date.

These clocks allow a research-operations lead to diagnose latency. A recently collected annual filing can contain a months-old effective fact. A source may publish quickly while a collector polls it weekly. An analyst may review yesterday's collection only after a mandate change. Each condition calls for a different repair.

## Move from record age to field risk

The second change is to classify each field on two axes:

- **Volatility:** how often the fact could change between observations.
- **Decision consequence:** how much a wrong or unknown value could distort identity, qualification, ranking or outreach.
- **Observability:** whether an authoritative source emits events, exposes periodic snapshots, or requires manual checking.
- **Jurisdiction:** which registry or reporting obligation governs the entity.
- **Evidence quality:** whether the source establishes the fact directly or only supplies an estimate or proxy.

This matters because official calendars differ. Australian Business Number holders are told to update details within **28 days** of learning of changes, and online changes take effect immediately (abr.gov.au). Canadian federal corporations must report a registered-office change within **15 days** (ised-isde.canada.ca), while New York corporations file a biennial statement every **two years** (dos.ny.gov). These are reporting requirements, not promises that every record is complete or promptly corrected.

## Put event triggers ahead of blanket schedules

The third change is architectural. Push-capable sources should create review work when a relevant event occurs. Periodic schedules should cover sources without dependable change signals.

- **New filing:** re-check status, name, address, officers or ownership fields affected by the filing.
- **Control event:** re-research ownership and all screening criteria dependent on owner type.
- **Domain redirect or certificate change:** verify identity before reusing website-derived facts.
- **Service-page change:** re-check vertical, offering and geography assertions supported by that page.
- **Bounced contact:** mark the address invalid and trigger contact re-research.
- **Mandate change:** evaluate the new question across every in-scope target, regardless of record age.
- **Pre-handoff gate:** refresh high-consequence fields before another team acts on the record.

Companies House says its stream can keep an imported snapshot current (developer-specs.company-information.service.gov.uk), and its metadata records when the event was published (developer-specs.company-information.service.gov.uk). The SEC submissions API typically processes in less than **one second**, whereas its XBRL interfaces typically process in under **one minute** and can take longer at peaks (sec.gov). Those are availability properties after dissemination. They do not say when a private company changed, whether it had to report the change, or whether the reported field answers the acquisition question.

### Preserve stale, unknown, contradicted and revalidated

A binary current-or-expired flag destroys useful information. The minimum field-state model is:

- **Current:** within policy age, supported by retained evidence, with no unresolved conflict.
- **Stale:** older than its maximum research age or due for a decision-gate review.
- **Unknown:** no supported answer, including a null or exhausted public-source search.
- **Contradicted:** two credible sources support incompatible values or dates.
- **Revalidated:** checked again, with the prior value still supported and new review evidence retained.
- **Superseded:** replaced by a newer supported value, while history remains queryable.

The EU guidance says data are consistent when they contain no contradictions (op.europa.eu), and the UK framework also treats non-contradiction as a distinct quality property (gov.uk). Revalidation should therefore add a review event, not silently rewrite the old collection date. W3C provenance models derivation as the transformation of one entity into another (w3.org), while SQL Server temporal tables can automatically store a row's previous version (learn.microsoft.com). Both patterns support auditable change rather than destructive replacement.

## Implementation Considerations and Process Changes

### Set decision gates before setting day counts

A private equity research refresh cadence should begin with the action that will consume the answer. A practical sequence is:

1. **Universe construction:** tolerate older low-consequence descriptors if the record is clearly marked and will not drive outreach.
2. **Screening:** require fresh identity, legal status, domain, ownership and core fit fields.
3. **Ranking:** refresh any field whose value materially changes the score or disqualification result.
4. **Handoff:** require no unresolved stale high-risk field unless the recipient explicitly accepts it.
5. **Approved outreach:** re-check company identity, domain, contact route and mandate exclusions.
6. **Diligence escalation:** replace proxy estimates with transaction-appropriate primary evidence.

The 2025 International Private Equity and Venture Capital Valuation Guidelines frame information around what was known or knowable at the measurement date (privateequityvaluation.com) and recognize that underlying performance information can be months in arrears (privateequityvaluation.com). Although valuation and origination are different workflows, the transferable lesson is to state the decision date and the information date separately.

### Compare operating models without confusing them

*Table 1* compares operating modes. It is not a provider ranking. A firm can combine all three, and the right design depends on list size, source coverage and internal review capacity.

| Operating model | Best use | Freshness control | Evidence and limitation |
|---|---|---|---|
| **Direct source feeds** | High-volume fields exposed by registries or filing systems | Consume events, retain source publication time, reconcile to a snapshot | Companies House can resume a stream from a snapshot timepoint (developer-specs.company-information.service.gov.uk). A feed covers only reported fields and jurisdictions. |
| **Internal research operations** | Mandate-specific judgments, estimates and contradiction resolution | Field SLAs, queues, analyst review and retained history | The UK framework recommends prioritizing fields or tables with the greatest user impact (gov.uk). Capacity and consistency must be measured locally. |
| **Docket, Triage, Scout and Audit** | Sourced target research delivered self-serve or managed | Store source, excerpt and collection date, then apply the buyer's field policy and decision gates | Docket says it retains the source page, supporting excerpt and collection date. That record still needs effective, publication, review and next-review semantics. |

The comparison shows why automation does not eliminate research judgment. A source feed reduces collection latency. It cannot decide whether a registered office is an operating location, whether a legal owner satisfies a mandate's ownership definition, or whether a revised service page changes sector fit. Docket is included because it directly supplies sourced target research, but no claim is made that one operating model guarantees fresher or more accurate answers.

## Store facts as observations, not mutable cells

The following pseudocode keeps the four clocks and the policy result distinct. `max_age_days` and `minutes_per_field` are organization-entered settings, not benchmarks.

```
SELECT
  target_id,
  field_name,
  field_value,
  as_of,
  source_published_at,
  collected_at,
  reviewed_at,
  next_review_at,
  CASE
    WHEN field_value IS NULL THEN 'unknown'
    WHEN contradiction_count > 0 THEN 'contradicted'
    WHEN CURRENT_TIMESTAMP >= next_review_at THEN 'stale'
    WHEN reviewed_at > prior_reviewed_at THEN 'revalidated'
    ELSE 'current'
  END AS freshness_state
FROM target_fact_observations;
```

The implementation should also store `source_url`, a retained evidence object, extraction or quotation, reviewer identity, jurisdiction, confidence and superseded observation identifier. W3C PROV defines generation time as when an entity was completely created and available for use (w3.org), and NIST's Open

Security Controls Assessment Language requires last-modified metadata to be assigned when content changes (pages.nist.gov). These are useful provenance patterns, not acquisition-specific mandates.

## Freshness Policy Matrix

The matrix below is a starting policy for a PE origination team. Every maximum age is an **illustrative user-entered setting**, not a published benchmark. A team should shorten or lengthen it based on mandate sensitivity, observable events, source lag and cost of error.

*Table 2* gives each field its own decision use, source, trigger and stale action.

| Field | Decision use and volatility | Authoritative source and known latency | Illustrative maximum research age | Event trigger and fallback | Action when stale |
|---|---|---|---|---|---|
| **Legal status and legal name** | Identity and eligibility, low frequency but very high consequence | Governing registry. ASIC open register data are updated weekly (asic.gov.au). | **30 days** | Filing/status event; fallback to registry search and retained filing | Block handoff or outreach until checked. |
| **Ownership or control** | Owner-type screen and conflict check, medium volatility and high consequence | Registry control filing where available. GLEIF Level 2 addresses who owns whom (gleif.org). | **30 days** | Control filing, merger announcement, mandate change; fallback to filings and company evidence | Re-research before scoring or approach. Preserve conflicts. |
| **Company domain** | Identity resolution and source routing, medium volatility and high consequence | RDAP plus company-controlled site. RDAP “last changed” concerns the registry object, not site content (datatracker.ietf.org). | **14 days** | Redirect, DNS or certificate change, site failure; fallback to registry and filings | Quarantine website-derived facts until identity is reconfirmed. |
| **Headcount** | Size screen and capacity estimate, high volatility and medium consequence | Company disclosures and dated job or staffing evidence. BLS says target-level employee count is confidential (bls.gov). | **90 days** | New disclosure, hiring change, transaction update; fallback to multiple dated proxies | Re-estimate and show method, range and date. |
| **Services and vertical** | Fit and thesis mapping, medium volatility and medium to high consequence | Company site, filings, catalogues and dated announcements | **90 days** | Service-page change, launch, acquisition, mandate taxonomy change; fallback to archived retained evidence | Refresh affected classifications, not the entire record. |

| Field | Decision use and volatility | Authoritative source and known latency | Illustrative maximum research age | Event trigger and fallback | Action when stale |
|---|---|---|---|---|---|
| **Geography** | Territory fit and coverage, medium volatility and medium consequence | Registry plus operating-location evidence. Companies House defines its address as the registered office (developer-specs.company-information.service.gov.uk). | **90 days** | Address filing, location-page change, closure or expansion notice; fallback to company sources | Distinguish registered office from operating footprint. |
| **Revenue estimate** | Size screening and prioritization, high uncertainty and high consequence | Current filed financials where available, otherwise dated triangulation. SEC warns grouped facts can have different reporting periods (sec.gov). | **90 days** | New accounts, financing, acquisition or direct disclosure; fallback to range with assumptions | Recalculate before a size-based decision and retain the old estimate. |
| **Contact details** | Approved outreach, high volatility and high consequence at action time | Company-controlled channel and delivery response. Mail status X.1.1 means the mailbox does not exist (datatracker.ietf.org). | **7 days before use** | Bounce, role change, opt-out or domain change; fallback to switchboard or company form | Do not reuse the failed route; re-research and preserve suppression controls. |

The matrix intentionally gives identity, ownership and contact routes shorter windows than descriptive fields. It also ties each stale state to an action. A stale service classification may remain usable for market mapping with a warning; a stale legal identity should stop an outbound handoff. Jurisdiction remains decisive: Canadian federal corporations have a **15-day** change filing rule for individuals with significant control (ised-isde.canada.ca), while California uses an event-triggered filing when a service agent's name or address changes (sos.ca.gov).

# Data Analysis and Evidence

## Source clocks do not imply fact currency

Official sources demonstrate why one cadence cannot fit all fields. GLEIF's Golden Copy and delta files update **three times daily** (gleif.org), yet legal entity reference data are reviewed and revalidated at least **annually** (gleif.org). New Zealand Business Number bulk files refresh **monthly**, while its API provides real-time access to newer data (nzbn.govt.nz). These sources can distribute a record promptly even when its underlying fact awaits the next entity report or validation.

AWS Glue illustrates explicit age rules. Its DataFreshness check compares current time with a selected date column, requires an hour or day expression, and fails null timestamps by default (docs.aws.amazon.com). This is an implementation analogy, not evidence that a particular number of hours suits acquisition research. Its separate FileFreshness rule uses a file's modification time (docs.aws.amazon.com), reinforcing that storage freshness can differ from fact freshness.

### Measure freshness debt and review load

The core portfolio metric is:

```
freshness_debt = stale_high_risk_fields / all_high_risk_fields
```

It should be paired with **p50 age** and **p90 age** for every material field, plus counts of unknown and contradicted answers. A median can look healthy while an old tail remains. PostgreSQL's continuous-percentile aggregate calculates a value corresponding to a chosen fraction (postgresql.org), and Prometheus documents a 90th-percentile query pattern (prometheus.io). Either implementation must define which timestamp supplies the age.

Expected review load is transparent arithmetic:

```
records x fields_due_per_period x minutes_per_field
```

Every input is a **user-entered assumption**. For example, 400 records multiplied by 2.5 fields due per month multiplied by 6 minutes per field equals **6,000 minutes**, or **100 analyst hours per month**. This is a capacity scenario, not an observed industry benchmark. Teams should replace all three inputs with their own measured values, then compare forecast load with completed review time.

## Case Studies and Real-World Examples

### Fictional batch handoff (Hypothetical Example)

Consider a fictional batch of **200 acquisition targets** with **four high-risk fields** each: legal status, ownership, domain and contact route. The policy therefore governs **800 high-risk field observations**. On handoff morning, 72 observations are beyond policy age, 18 are unknown and 10 are contradicted. Freshness debt is **72 / 800 = 9%**. Unknown and contradicted observations are reported separately because neither is simply stale.

*Table 3* shows the fictional queue. Every time and quantity is a user-entered example, not a benchmark.

| Field | Observations | Stale | Unknown | Contradicted | p50 / p90 age | Handoff action |
|---|---|---|---|---|---|---|
| **Legal status** | 200 | 8 | 1 | 0 | 9 / 28 days | Event check, then direct registry review for the 8 stale records. |
| **Ownership** | 200 | 20 | 9 | 10 | 18 / 61 days | Re-research 20 stale and resolve 10 conflicts before scoring. |
| **Domain** | 200 | 16 | 2 | 0 | 5 / 21 days | Revalidate redirects and identity before trusting website evidence. |
| **Contact route** | 200 | 28 | 6 | 0 | 3 / 13 days | Re-check immediately before approved outreach. |

The table directs work only to affected high-risk fields, not all fields on all 200 targets. If the team assumes **6 minutes per due field**, the 72 stale observations imply **432 minutes**, or **7.2 hours**, before accounting for contradiction resolution. A blanket rebuild of all 800 observations at the same assumed rate would require **80 hours**. The comparison is arithmetic from fictional assumptions, not evidence that selective refresh will always save that amount.

The queue also responds to events. A newly accepted filing can trigger a legal-status check; Companies House describes immediate filing alerts after acceptance (gov.uk). For public-company events within Form 8-K scope, the SEC generally requires filing within **four business days** and includes an item for change in control (sec.gov). Those event paths reduce detection delay for covered facts, but they do not extend to every private target or every operational change.

## Implications and Future Directions

**Private equity target monitoring** should evolve from list-level recrawls to observation-level policies. Change tracking is a useful engineering pattern: SQL Server records the version in which a row last changed (learn.microsoft.com), and Apache Iceberg uses an increasing sequence number to order table changes (iceberg.apache.org). Acquisition research can apply the same idea by storing immutable observations and emitting field-specific review tasks.

The next improvement is feedback from actual workflow outcomes. Teams can measure:

- **Detection lag:** event effective time to source publication, then publication to collection.
- **Review lag:** collection to analyst decision.
- **Revalidation yield:** proportion of scheduled reviews that change the accepted value.
- **Contradiction rate:** share of reviewed fields with unresolved credible conflicts.
- **Unknown closure rate:** share of unknowns later resolved by a new source.
- **Decision-gate debt:** high-risk stale fields remaining at screening, handoff or outreach.
- **Source coverage:** share of fields with push events, snapshots only, or manual checks.

These measures allow the SLA to learn. A field with low change yield and strong event coverage may justify a longer scheduled interval. A field with repeated pre-outreach failures may need a shorter action-time check. DataCite’s versioning guidance offers a transferable provenance pattern by preserving metadata changes in an activities endpoint (support.datacite.org). The objective is not maximum refresh volume. It is the smallest observable review load that keeps decision-critical uncertainty explicit.

Future systems should also separate source confidence from freshness. A freshly fetched website assertion may still be ambiguous. A year-old filed historical fact may be exactly right for a period-specific question. The research interface should display both quality dimensions rather than collapsing them into a colored badge.

# Frequently Asked Questions (FAQs)

## How often should acquisition target research be updated?

There is no defensible universal interval. Use event triggers for reportable changes, set field-specific maximum ages for coverage gaps, and force checks at screening, handoff and outreach. High-consequence identity, ownership, domain and contact fields generally deserve shorter locally chosen windows than descriptive fields. W3C's formulation is the right principle: currentness is the right age for a specific context of use (w3.org).

## What is a practical private equity research refresh cadence?

A practical cadence is two-layered: continuous or daily event ingestion where sources support it, plus periodic field queues where they do not. GLEIF receives daily LEI and reference-data reporting (gleif.org), while Companies House daily accounts files contain records registered the preceding day (gov.uk). The research team still sets a review cadence based on field risk and its observed backlog.

## What belongs in an acquisition target re-research SLA?

The SLA should state the field, decision use, source hierarchy, source latency, maximum age, triggering events, fallback evidence, stale action, jurisdiction, owner and escalation path. It should also define stale, unknown, contradicted, revalidated and superseded. FINRA's guidance in a different diligence setting calls for a process tailored to unique facts and circumstances (finra.org). The same tailoring principle is appropriate here.

## What is the right due diligence data refresh frequency?

Origination research and transaction due diligence should not share one threshold. As a target progresses, direct company information and transaction materials should replace public proxies. At every gate, refresh facts that could change the decision and retain the evidence used at that time. A current collector timestamp alone is never sufficient.

## When should company research be refreshed automatically?

The decision about **when to refresh company research** automatically starts with source semantics. Automate refresh when the source emits a machine-readable event that maps clearly to a field, when identity matching is reliable, and when the result can be reviewed or safely accepted under policy. Schedule or manually review ambiguous website changes, ownership inferences, revenue estimates and contradictions. Google Cloud's freshness template checks whether the latest timestamp exceeds an interval (cloud.google.com), but the organization must decide which timestamp and interval matter.

## What are the core M&A target research best practices for freshness?

- **Keep the clocks separate.** Store effective, publication, collection and review times.
- **Apply field-level policy.** Do not expire a whole record because one field is due.
- **Use events first.** Let filings, redirects, bounces and mandate changes create targeted work.
- **Retain evidence.** Preserve sources, excerpts, prior values and review decisions.

- **Show uncertainty.** Keep unknown and contradicted states visible.
- **Measure debt.** Report stale high-risk share plus p50 and p90 age by field.
- **Recheck before action.** Screening, handoff and outreach are independent freshness gates.

## Conclusion

Acquisition target research update frequency is not a choice between real time, monthly or quarterly. It is a control system that matches the age of each fact to volatility, consequence, source behavior and the next decision. The strongest policy separates `as_of`, `source_published_at`, `collected_at` and `reviewed_at`; it uses `next_review_at` only as a fallback when events have not already created work.

For PE research operations, the field matrix is the operative SLA. Legal status, ownership, domain and contact routes can stop a handoff when stale. Headcount, services, geography and revenue estimates can follow different evidence paths, with explicit ranges or unknowns where public sources cannot establish a current value. A stale answer is a review instruction, not an automatic declaration that the answer is false.

The management view should remain compact: freshness debt for high-risk fields, p50 and p90 age by field, unknown and contradiction counts, detection lag, review lag and revalidation yield. Capacity follows from records multiplied by due fields multiplied by a locally measured minutes-per-field input. This makes the cost of a tighter policy visible without presenting an invented industry norm.

No schedule guarantees correctness. Freshness controls age; retained evidence, source hierarchy, contradiction handling and analyst review control verifiability. The decision rule is simple: reuse when the field remains within policy and no trigger has fired, refresh automatically when an authoritative event maps cleanly to the field, and re-research when consequence, ambiguity or the next workflow gate demands judgment.

## Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.


1. https://www.sec.gov/search-filings/edgar-application-programming-interfaces
2. https://developer-specs.company-information.service.gov.uk/streaming-api/guides/overview
3. https://www.gov.uk/guidance/companies-house-data-products
4. https://www.bls.gov/cew/overview.htm
5. https://www.bls.gov/cew/questions-and-answers.htm
6. https://op.europa.eu/webpub/op/data-quality-guidelines/en
7. https://www.gov.uk/government/publications/the-government-data-quality-framework/the-government-data-quality-framework
8. https://www.w3.org/TR/vocab-dqv
9. https://icis.corp.delaware.gov/ecorp/FieldDesc.aspx
10. https://www.dublincore.org/resources/userguide/creating_metadata
11. https://www.abr.gov.au/business-super-funds-charities/updating-or-cancelling-your-abn/update-your-abn-details
12. https://ised-isde.canada.ca/site/corporations-canada/en/business-corporations/next-steps-following-incorporation-your-business

13. https://dos.ny.gov/biennial-statements-business-corporations-and-limited-liability-companies
14. https://www.w3.org/TR/prov-o
15. https://learn.microsoft.com/en-us/sql/relational-databases/tables/temporal/overview?view=sql-server-ver17
16. https://www.privateequityvaluation.com/Portals/0/Documents/Guidelines/2025%20IPEV%20Valuation%20Guidelines.pdf
17. https://www.gov.uk/government/publications/the-government-data-quality-framework/the-government-data-quality-framework-guidance
18. https://pages.nist.gov/OSCAL/learn/concepts/layer/overview
19. https://www.asic.gov.au/online-services/search-asic-registers/datagovau
20. https://www.gleif.org/en/lei-data/access-and-use-lei-data
21. https://datatracker.ietf.org/doc/rfc9083
22. https://developer-specs.company-information.service.gov.uk/companies-house-public-data-api/resources/companyprofile?v=latest
23. https://datatracker.ietf.org/doc/rfc3463
24. https://ised-isde.canada.ca/site/corporations-canada/en/individuals-significant-control/individuals-significant-control-file-your-information
25. https://www.sos.ca.gov/business-programs/business-entities/faqs
26. https://www.gleif.org/en/newsroom/blog/the-lei-in-numbers-active-lei-population-surpasses-3-million-in-q1-2026
27. https://www.nzbn.govt.nz/using-the-nzbn/nzbn-services/bulk-data
28. https://docs.aws.amazon.com/glue/latest/dg/dqdl-rule-types-DataFreshness.html
29. https://docs.aws.amazon.com/glue/latest/dg/dqdl-rule-types-FileFreshness.html
30. https://www.postgresql.org/docs/17/functions-aggregate.html
31. https://prometheus.io/docs/prometheus/3.9/querying/functions
32. https://www.gov.uk/guidance/searching-the-companies-house-register
33. https://www.sec.gov/divisions/corpfin/forms/8-k.htm
34. https://learn.microsoft.com/en-us/sql/relational-databases/track-changes/work-with-change-tracking-sql-server?view=sql-server-ver17
35. https://iceberg.apache.org/spec?h=changing
36. https://support.datacite.org/docs/versioning
37. https://www.finra.org/sites/default/files/2023-05/regulatory-notice-23-08.pdf
38. https://docs.cloud.google.com/dataplex/docs/system-rule-templates-list

THE PUBLISHER

# About Docket

Docket provides deal-origination research software and managed research for private equity firms. We help investment teams investigate acquisition targets using structured screening criteria, retained sources and reviewable company evidence. Teams can work through a self-serve platform or use managed research, depending on how they want research delivered.

### Research against a defined mandate

A useful target list needs more than company names. Docket focuses on the questions that determine whether a company fits an investment mandate, including the evidence needed to support or qualify each answer. Our research approach makes the connection between screening criteria, source material and conclusions visible to the team reviewing the work.

### Triage, Scout and Audit

Docket's three named research agents perform complementary tasks. Triage resolves company identity and screens fit. Scout collects sourced answers against the mandate. Audit checks retained evidence, addresses contradictions and leaves unsupported answers visibly unresolved. This structure helps reviewers distinguish established findings from missing information and questions requiring further investigation.

### Evidence that supports investment-team judgment

Our research library covers market mapping, screening criteria, private-company data, succession and ownership, source evaluation and evidence standards. These resources explain the methods and limitations behind origination research. Findings support a team's judgment; they do not establish that a company is for sale or guarantee a transaction or investment outcome.

### Work with Docket

Visit Docket to explore the platform, managed research and the current contact path. Read about screening criteria, evidence standards, AI research agents and market mapping.

Public examples are illustrative unless explicitly identified otherwise. Research preparation and authorized outreach are separate activities; confidential target lists and customer outcomes should never be inferred from an educational example.

**Website:** https://docket.capital

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.