Docket

DOCKET / RESEARCH & PRACTICAL GUIDANCE

Job Postings vs. Headcount Growth: A PE Evidence Protocol

> 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 · 2026-10-17 · job postings vs headcount growth · job postings data · private equity due diligence · hiring signals · headcount analysis · job posting analytics · replacement hiring · employer resolution

Original article: https://docket.capital/articles/job-postings-vs-headcount-growth



In this article

1. Executive Summary
2. Introduction and Background
3. Key Changes
4. Implementation Considerations and Process Changes
5. Skills and Functions as Advertised Requirements
6. Comparing Posting Activity over Time
7. Corroboration and Headcount Verification
8. Data Analysis and Evidence
9. Case Studies and Real-World Examples
10. Implications and Future Directions
11. Frequently Asked Questions (FAQs)
12. Conclusion

Executive Summary

Job postings and headcount growth measure different things. A verified posting supports a dated statement that an identified employer advertised a role with particular requirements. It does not establish that someone joined, that the workforce expanded, or that the business grew. Indeed explicitly says its postings are not a precise count of available jobs and describes replacement activity even without employment growth. (hiringlab.indeed.com)

The distinction matters in private equity (PE) origination because an advertised function can justify further research without satisfying a company-size or capability screen. Official statistics separate **openings**, **hires**, and **employment**: the US Bureau of Labor Statistics (BLS) measures openings on the last business day of the month and hires as payroll additions during the month. Different reference periods and objects prevent a posting total from serving as a workforce reconciliation. (www.bls.gov) The Organisation for Economic Co-operation and Development (OECD) likewise concludes that the Lightcast data it evaluated better reflect vacancy distributions than absolute levels. (www.oecd.org)

This report proposes a **posting-to-claim evidence protocol**: preserve the advertisement and collection history, resolve the employer, retain duplicate relationships, classify the advertised requirements, and assign a claim status. Provider definitions must travel with the data. Lightcast documents a **60-day** cross-source duplicate-check window. (kb.lightcast.io) TheirStack describes ignoring same-company, same-title reposts within **30 days**. (theirstack.com) SignalsAPI describes deduplication by company, title, and country. (jobs.signalsapi.com) These are provider descriptions of different counting rules, not independently established accuracy rankings.

The decision is therefore to **accept a narrow observation, seek corroboration, preserve an unknown, or reject an unsupported inference**. Geographic eligibility for remote work does not establish an office, and disappearance from a feed does not establish a successful hire. Google's documentation distinguishes remote work from attendance at a physical office; research using linked application and vacancy data also distinguishes posting duration from vacancy duration. (developers.google.com) (www.nber.org) The worked

example below uses explicitly fictional inputs to calculate duplicate share, employer-resolution yield, stale share, and unknown share. It deliberately produces no headcount-growth estimate. A posting should change the next research question only to the extent that its identity, timing, and literal content support that question.

Introduction and Background

Do job postings indicate company growth? They can indicate advertised recruitment activity within the sources observed. Growth becomes a separate proposition requiring evidence about the relevant workforce, business perimeter, and period. The Australian Bureau of Statistics (ABS) explicitly distinguishes vacancy counts from advertisement counts, including cases where a vacancy is advertised multiple times. (www.abs.gov.au) That distinction should remain visible when an acquisition screen reduces detailed research to a simple answer.

The commercial data category includes **job postings datasets**, databases, application programming interfaces (APIs), and job posting analytics. PandorLabs, for example, describes collection from company career pages and major boards. (www.pandorlabs.com) The product description establishes what the vendor says it collects. An analyst still needs to determine whether an individual record belongs to the target and supports the proposed claim.

The useful question is narrower than whether a target is “growing”: **What did this employer publicly advertise, when was it observed, and what does that evidence permit the researcher to say?** A specialist engineering advertisement may support a research lead about a technical function. It should not automatically satisfy a screen requiring demonstrated delivery capability.

The methodology pages were accessed on October 17, 2026; historical findings retain their original periods. Docket's perspective in this report is that of a target-research publisher. The scope is the evidentiary treatment of hiring signals for acquisition target screening. It does not rank acquisition candidates or infer sale intent.

The protocol begins with the source record and ends with a reviewable claim. Its recommendations are proposed research controls, rather than claims of predictive performance. Official definitions supply the boundaries: a job advertisement is a listing about an opening, while official employment statistics can count jobs differently from unique people. (schema.org) (www.ons.gov.uk) Keeping those boundaries intact makes a posting useful without forcing it to answer an employee-count question.

Key Changes

Change the unit before changing the conclusion

The first change is to name the object being counted. A **raw ad** is an observed source item. A **deduplicated posting** is a cluster representative under a specified rule. A **requisition** is an employer's hiring record when its identity is actually available. An **opening** is a position meeting the applicable definition. A **hire** is an addition to the workforce or payroll under the measurement used. BLS defines hires through payroll additions, while Schema.org describes a JobPosting as a listing about an opening. (www.bls.gov) (schema.org)

These objects do not have a universal conversion factor. Schema.org's totalJobOpenings property allows publishers to identify the positions represented by a posting, but says not to use the property when that number is unclear or unknown. (schema.org) Missing multiplicity should remain unknown.

A public post need not even identify a standard requisition. Greenhouse documents prospect posts whose internal job identifier is null. (docs.greenhouse.io) An analyst should retain that distinction rather than assigning every advertisement an invented requisition number.

Separate replacement hiring from net headcount growth

Replacement hiring replenishes departures; **net headcount growth** concerns a change in the employee stock within a defined perimeter. Indeed explicitly identifies job changing without employment growth. (hiringlab.indeed.com) Consequently, more advertised roles can coexist with stable employment.

A useful workforce bridge is beginning employees plus starters, minus leavers, with separately disclosed changes in the entity perimeter and measurement basis. BLS notes that hires minus separations can derive a net employment-change measure, while explaining differences from its payroll employment series. (www.bls.gov) A company reconciliation should therefore specify its own dates and population rather than copy a national-statistics formula without checking scope.

Preserve conflicting count labels

Lightcast's total-postings page explicitly defines total as the count before deduplication, although its opening description also uses “unduplicated.” Its separate unique-postings definition consistently identifies the count after deduplication. Use the explicit definitions and retain the wording discrepancy in the source notes. (kb.lightcast.io)

Even “employment” needs a unit. BLS payroll statistics count someone working at different establishments at each establishment; ABS distinguishes employed people from total jobs, which include filled and vacant jobs. (www.bls.gov) (www.abs.gov.au)

For a PE screen, prohibit the following shortcuts:

- **Posting count to employees:** No conversion without workforce evidence.
- **Advertised roles to net hires:** Departures and completed starts remain unobserved.
- **Role disappearance to successful hire:** Removal has no verified outcome by itself.
- **Hiring activity to revenue:** Require financial evidence for a revenue claim.
- **Recruitment to operating health:** Keep that broader judgment outside the posting observation.
- **Advertising to sale intent:** No acquisition-readiness conclusion follows.

These are recommended inference boundaries. They let a researcher accept a literal observation while rejecting the larger statement proposed for a screening field.

Implementation Considerations and Process Changes

Capture provenance before enrichment

Store both the page and the record lineage. Google's datePosted means the employer's original posting date; it is different from the analyst's collection date. (developers.google.com) TheirStack similarly distinguishes where it first saw a job from where that job was published. (theirstack.com)

A practical **posting evidence worksheet** should contain:

- **Source record:** Raw ID, URL, source class, and retained text.
- **Hiring identifiers:** Public post ID and requisition ID, with unknowns preserved.
- **Time evidence:** Stated posting date, first seen, last seen, and collection date.
- **Entity evidence:** Display name, proposed employer, legal-entity crosswalk, and confidence.
- **Role evidence:** Literal title, body, stated location, and remote conditions.
- **Duplicate lineage:** Cluster ID, retained representative, rule version, and collisions.
- **Extracted attributes:** Skill, function, taxonomy release, and supporting sentence.
- **Review outcome:** Corroboration, conflicting evidence, claim status, and next source.

Do not assume that an API's similarly named identifiers represent the same object. Greenhouse distinguishes a public job-post identifier from the underlying internal job identifier. (docs.greenhouse.io)

Resolve the employer, then deduplicate

Employer resolution should precede company-level interpretation. Google expressly accommodates a staffing agency acting for an anonymous employer. (developers.google.com) BLS assigns temporary help workers, contractors, and consultants to their employer of record for its relevant statistics. (www.bls.gov) Neither a recruiting brand nor a worksite name alone resolves the employee relationship.

Separate brand, franchisee, operating subsidiary, parent, and facility in the worksheet. Census defines an establishment as a physical economic location and aggregates firm data across establishments under common ownership. (www.census.gov) GLEIF's parent data concern accounting-consolidating parents, which help map a group but do not identify the employer of a particular worker. (www.gleif.org)

Table 1 compares evidence inputs and the research layer. Provider entries describe documented methods, without asserting equivalent coverage or accuracy.

Input or research layer	Documented basis	Permitted role in the protocol
Employer ATS, illustrated by Greenhouse	Published jobs, offices, and departments are exposed through its Job Board API. (docs.greenhouse.io)	Locate the employer-linked source and retain identifiers.
Lightcast	Unique postings are counted after deduplication. (kb.lightcast.io)	Analyze the provider-defined posting unit.
SignalsAPI	Company, title, and country define its described deduplication unit. (jobs.signalsapi.com)	Preserve geographic aggregation rules.

Input or research layer	Documented basis	Permitted role in the protocol
TheirStack	The provider describes continuous crawling. (theirstack.com)	Record observation timing separately from employer publication.
PandorLabs	Clustering uses company, normalized title, location, and description similarity. (www.pandorlabs.com)	Retain cluster lineage and test ambiguous matches.
Docket, publisher perspective	Target-research interpretation under the protocol proposed here.	Apply the same acceptance, corroboration, and unknown rules to each input.

The table compares roles in an evidence process, not interchangeable product specifications. Its counting differences are material. Repeated descriptions can belong to one job even when their wording varies, as a deduplication research paper explains. (arxiv.org) Conversely, similar titles should not automatically merge distinct requisitions.

Use a transparent funnel: **raw ads → source-normalized ads → duplicate clusters → unique observed postings → verified employer-linked postings**. Keep discarded duplicates accessible. Reposts, evergreen recruiting, location variants, and applicant tracking system (ATS) migrations should trigger review of continuity, not an automatic new-opening count.

Skills and Functions as Advertised Requirements

A posting is useful for capability screening when the research preserves what the text actually requests. **Requested skill**, **preferred skill**, **role responsibility**, and **shared employer language** should be separate attributes. Greenhouse can return content assembled as a concatenated string, which makes an undifferentiated text field an insufficient guide to which passage belongs to the role. (docs.greenhouse.io)

Use a frozen taxonomy to make extraction reproducible. European Skills, Competences, Qualifications and Occupations (ESCO) uses concept uniform resource identifiers (URIs) as keys for joining its datasets. (esco.ec.europa.eu) The Occupational Information Network (O*NET) separates skills and knowledge within its model of worker requirements. (www.onetcenter.org) Neither classification turns an advertised requirement into proof that the employer currently possesses that capability; that is an analytical limit of the evidence.

Recommended extraction controls are:

- **Keep the sentence:** Store the precise passage behind every assigned skill.
- **Preserve modality:** Separate required, preferred, optional, and training-provided language.
- **Respect negation:** A statement that experience is unnecessary should not become a requirement.
- **Separate functions:** Classify the role's work independently from the employer's industry.
- **Exclude shared text:** Review mission statements, benefits, and equal-opportunity language separately.
- **Freeze identifiers:** Retain taxonomy concept IDs alongside displayed labels.
- **Retain ambiguity:** Permit multiple plausible classifications pending review.
- **Review consequential labels:** Manually inspect attributes that decide mandate eligibility.

Versioning matters because O*NET describes quarterly database updates and maintains release archives. (www.onetcenter.org) A classification change should create a methodology flag. Otherwise an apparent increase in a function's posting mix might reflect revised labels.

ESCO also distinguishes language-dependent labels from language-independent concept relationships. (esco.ec.europa.eu) For multilingual data, preserve the original phrase and translation separately from the assigned concept. A reviewer should be able to see whether a technical term came from the source or from an enrichment step.

The accepted wording should stay close to the evidence: “The employer advertised a role requiring the specified skill.” A stronger statement, such as “The company delivers this service,” needs separate operational evidence. O*NET's structured model describes jobs and workers; it does not certify an individual employer's service portfolio. (www.onetcenter.org) Function and skill mixes can guide follow-up research, but the population remains **observed advertisements with classifiable text**.

Comparing Posting Activity over Time

A trend requires a stable measurement contract. Decide whether the series measures an **active stock** or a **new-posting flow**. Indeed counts new postings when they first become visible; Lightcast's UK glossary describes counting an active posting once in each month it remains active. (hiringlab.indeed.com) (docs.lightcast.io) Summing those monthly active counts would therefore not identify unique jobs over the whole interval.

Freeze the employer perimeter, sources, query, deduplication, location handling, and time window before comparing periods. Lightcast documents that improved expiration methodology can change historical active-posting trend lines. (docs.lightcast.io) Indeed also warns that platform presentation changes can resemble demand shifts. (www.indeed.com)

The comparison log should record:

- **Source continuity:** Which career pages, boards, and endpoints were observed in both periods.
- **Collection health:** Successful checks, failed checks, and unavailable intervals.
- **Backfills:** Older records newly discovered after collection or coverage changes.
- **Identity changes:** Renames, acquisitions, franchise boundaries, and subsidiary remapping.
- **Rule changes:** Duplicate windows, matching thresholds, and taxonomy releases.
- **Window alignment:** Identical exposure periods and a consistent stock or flow definition.

First observed is not necessarily newly published. Last observed is not necessarily the employer's closing date. A missing collection interval should remain a gap, rather than become a zero.

Retention limits also constrain interpretation. SignalsAPI says its market-page surface retains a rolling **30-day posting-date window**. (jobs.signalsapi.com) That surface alone cannot establish a longer history or a complete vacancy lifecycle. This is a documentation-based limitation of the observed surface, not a claim about every product the provider might offer.

For an optional within-company index, define the base-period posting measure as the denominator and multiply the ratio by a chosen scaling constant. Label the result an **observed-posting index**. Suspend comparisons when the denominator is zero or coverage changes cannot be reconciled. This is a proposed descriptive measure, not a headcount-growth model.

Cross-company league tables need stronger coverage evidence. The OECD's European validation warns against assuming equal coverage across countries, while Cedefop assesses representativeness at sectoral, occupational, and geographic levels. (www.oecd.org) (www26.cedefop.europa.eu) A provider's broad source list does not, by itself, establish comparable observation probability for two acquisition targets.

Official comparison series also have different smoothing. The UK Office for National Statistics (ONS) bases headline vacancy figures on **three-month moving averages**. (www.ons.gov.uk) Comparing an unsmoothed company's daily ad count with that headline series requires an explicit timing adjustment and still cannot establish causation.

Corroboration and Headcount Verification

How to verify company headcount growth begins with a workforce definition, not with the posting count. Specify legal entities, geography, employees versus contractors, full-time equivalent (FTE) versus people, and point-in-time versus average counts. Then seek workforce observations on a comparable basis. BLS's separate definitions of hires and separations illustrate why both additions and departures matter. (www.bls.gov)

For public companies, read the human-capital disclosure and its scope. The Securities and Exchange Commission (SEC) describes this disclosure through a materiality-based framework; its final rule retained the number of persons employed where material to understanding the business. (www.sec.gov) Do not assume that a filing supplies a standardized monthly workforce bridge.

For private-company due diligence, request dated employee registers and a reconciliation when authorized access becomes available. For public-source screening, retain dated company statements and explicit limitations. If comparable workforce evidence is unavailable, the answer to headcount growth remains **unknown**, even when recruitment activity is visible.

Table 2 maps observable fields to narrow claims and the evidence needed to go further.

Observed field	Safe claim	Unsupported inference	Next evidence to seek
Employer-linked role and identifier	This employer advertised the stated role when observed.	A completed hire or net workforce increase.	Dated workforce evidence and a hiring outcome.
Skill requirement in role text	The advertisement requested that skill.	The company currently delivers that capability.	Current service documentation, licenses where applicable, or staff evidence.
Remote applicant geography	Applicants may work from the stated area under the ad's conditions.	A physical office exists there. (developers.google.com)	Company location page and relevant facility records.

Observed field	Safe claim	Unsupported inference	Next evidence to seek
Missing position count	The advertised multiplicity is unknown.	The ad represents exactly one opening. (schema.org)	Employer-confirmed position count.
Salary field	Record its stated provenance and whether it is disclosed or estimated.	Confirmed payroll expense or approved total budget. (schema.org)	Employer compensation evidence and workforce scope.
Posting no longer observed	The source item disappeared from the observed surface.	The vacancy was filled or the role began. (www.nber.org)	Employer-confirmed outcome.

The matrix deliberately separates a text observation from a business conclusion. Unknown is a valid research result, not a reason to replace missing evidence with a ratio.

Salary illustrates the same rule. PandorLabs says its salary field remains null when compensation is undisclosed rather than estimated. (www.pandorlabs.com) Do not impute a salary, multiply by ads, and describe the result as a hiring budget.

Corroboration must also add independent information. A board copy of the same employer text confirms dissemination, not a separate hiring outcome. Preserve conflicting dates, entities, and wording. Accept a broader capability or location claim only when its separate evidence actually supports the broader proposition.

Data Analysis and Evidence

The empirical evidence supports measurement controls rather than a universal conversion from postings to growth. The OECD's English-speaking-country validation covers Australia, Canada, the United Kingdom, and the United States, with a main comparison period from **Q1 2015 to Q2 2022**. (www.oecd.org) Its conclusion favors distributions over absolute vacancy levels. (www.oecd.org) That finding concerns the studied dataset and jurisdictions; it is not a validation of company-level acquisition screening.

Employer identity was incomplete in that study's extract. For **January to June 2022**, employer-name missingness was **32% in Australia, 9% in Canada, 42% in the United Kingdom, and 19% in the United States**. These are historical study observations, not current provider coverage estimates. (www.oecd.org) They explain why employer-resolution yield deserves its own denominator.

Official survey designs differ too. BLS requires an opening that could start within **30 days**. (www.bls.gov) It measures employment using the pay period containing the **12th**. (www.bls.gov) ONS documentation describes questionnaires to approximately **6,100 businesses monthly**. (www.ons.gov.uk) Statistics Canada's **2023 guide** describes a quarterly sample of **100,000 locations** and recruitment outside the organization as part of its vacancy definition. (www150.statcan.gc.ca) These methodological differences make jurisdiction and reference period essential metadata.

Linked data do not erase the distinction between advertisements and outcomes. The NBER working paper *Application Flows* analyzes **125 million applications** linked to vacancies and employer-side clients on the technical-skills platform Dice.com. (www.nber.org) Its evidence that posting durations differ from vacancy durations is platform-specific, rather than a universal time-to-fill calibration. (www.nber.org)

Table 3 defines a small dashboard. Every input and calculation in its final column is **fictional illustrative arithmetic**, not an empirical finding or benchmark.

Metric	Explicit denominator and definition	Hypothetical Example calculation
Duplicate share	Excess raw records removed divided by raw records reviewed.	From 240 raw ads to 150 unique observed postings: 90 / 240 = **37.5%**.
Employer-resolution yield	Verified employer-linked postings divided by unique observed postings.	120 / 150 = **80%**.
Stale share	Postings failing the chosen recency rule divided by verified employer-linked postings.	18 / 120 = **15%**.
Unknown share	Claims remaining unresolved divided by all verified employer-linked claims reviewed.	24 / 120 = **20%**.
Advertised function mix	In-scope postings assigned a function divided by all in-scope postings in that analysis.	24 engineering postings / 96 classifiable postings = **25%**.

The denominator changes with the question. Duplicate share describes collection redundancy; employer-resolution yield describes attribution; unknown share describes unanswered claims. None estimates workforce expansion. The function-mix denominator excludes unclassifiable text, which must be disclosed rather than silently represented as a complete company workforce.

These metrics are proposed controls. Actual implementations should publish their exclusion rules, overlaps, and rule versions, and retain the counts needed to reproduce each result.

Case Studies and Real-World Examples

Posting-to-claim review (Hypothetical Example)

All company names, observations, dates, and counts in this example are fictional. Alder Process Services is an invented acquisition-screening record. The researcher observes 240 ads across a career page, an ATS feed, and boards during an illustrative review window. Source normalization preserves 240 records; clustering produces 150 unique observed postings. Employer review links 120 to the intended entity and leaves 30 without sufficient attribution.

The example uses the same arithmetic as Table 3. The 90 excess raw records represent duplicate observations. They remain in the retained record but are not treated as additional observed postings. The 30 employer-unresolved postings remain outside company-specific interpretation; exclusion from that interpretation does not erase their source history.

The reviewer then assesses candidate claims:

- **Accept:** An employer-linked advertisement requests a particular technician skill. Accept the dated advertising observation.
- **Corroborate:** A role names a new service function. Seek current delivery evidence before accepting a capability claim.

- **Stale:** An older board copy remains visible after the chosen freshness threshold. Keep the historical observation and withhold a current-opening statement.
- **Duplicate:** A syndicated copy repeats an identified employer posting. Link it to the cluster and preserve the copy's observation date.
- **Unresolved:** An agency advertisement names a broad industry and city but does not identify its client. Keep the end employer unknown.
- **Reject inference:** A posting disappears. Reject the proposed statement that an employee started, while retaining the disappearance observation.

The stale threshold is a fictional policy choice, not evidence that every older advertisement is closed. Likewise, a source removal is a collection event until the employer's outcome is known.

The final screen can say that Alder advertised specified roles and that selected capability questions require corroboration. It cannot say that the company added 120 employees, opened a facility, increased revenue, or became more likely to sell. Those claims were never observed.

The handoff should include an accepted sentence, its literal support, the employer link, observation dates, unresolved conflicts, and the next evidence needed. A reviewer can then accept the limited finding without accepting a chain of broader assumptions. This preserves a useful research lead while keeping the headcount-growth field unanswered.

Implications and Future Directions

For **job posting data for private equity screening**, the practical value is a better next question. Advertised functions can prioritize capability checks, and stated work locations can prioritize footprint checks. TheirStack explicitly notes that companies without published postings are absent from its posting data. (theirstack.com) Therefore, a silent feed should not be scored as a company with no hiring, no employees, or no relevant capability.

Procurement should test the decision the dataset must support. Ask providers for sample records, source lineage, duplicate-cluster explanations, timestamp definitions, historical revisions, and unresolved-employer treatment. Broad collection coverage and rapid refresh are different properties from accurate attribution or comparable employer observation.

Recommended acceptance tests are:

- **Attribution:** Can a reviewer reproduce the employer link?
- **Multiplicity:** Can the dataset distinguish ads from stated positions?
- **Continuity:** Can the same posting survive reposts and source changes without an invented new event?
- **Reproducibility:** Can the prior result be rerun under its original rules?
- **Unknown handling:** Does missing evidence stay visibly missing?
- **Claim fit:** Does the cited passage support the actual screening answer?

Benchmarking should match the analytical level. The OECD recommends comparing samples with other vacancy sources at an aggregation relevant to the use case. (www.oecd.org) National agreement would not, by itself, prove that a feed accurately tracks an individual small employer.

Future improvements should focus on retained evidence and measurable review outcomes: accepted employer links, reproducible clusters, stable source panels, and fewer unsupported conclusions. A claimed leading indicator or predictive score needs independent validation for its intended population and outcome. Until that evidence exists, the protocol should treat it as a proposition to test.

Frequently Asked Questions (FAQs)

Do job postings indicate company growth?

They indicate observed advertising activity when the source, employer, and timing are verified. Growth requires additional evidence. Replacement activity can occur without employment growth. (hiringlab.indeed.com)

Are job postings a growth indicator or a count of actual hires?

They are a potential research signal, not a count of actual hires. BLS defines hires as payroll additions, a different event from publication. (www.bls.gov) Validate the outcome before using the posting to support a workforce claim.

What belongs in headcount analysis for private equity due diligence?

Define the workforce perimeter, reconcile comparable employee observations, and account for starters, departures, and entity changes. Keep advertised hiring activity in a separate evidence field.

What makes a job postings dataset useful?

Useful job posting analytics preserve provenance, counting rules, collection windows, employer uncertainty, and reviewable text. A dataset's usefulness depends on the proposed claim and demonstrated coverage of the relevant employers.

How should analysts handle replacement hiring versus net headcount growth?

Keep replacement intent unknown unless the employer states it or workforce evidence resolves it. An advertisement alone cannot identify the balance between additions and departures.

Conclusion

Job postings versus headcount growth is a question about evidence boundaries. A posting can support a precise statement about an advertisement. It can also identify a capability, function, or location worth checking. It cannot supply the unobserved workforce movements needed to establish net employee growth.

The proposed protocol begins with provenance and employer identity, continues through transparent duplicate handling and frozen classifications, and ends with a claim status. It keeps the original text, timestamps, assumptions, and conflicting evidence available for review. That record makes it possible to accept a useful narrow observation without endorsing a broader unsupported interpretation.

For an origination team, the central discipline is to keep the research question attached to the evidence unit. Count observed postings when studying advertising. Seek operational evidence when assessing capability or footprint. Seek comparable workforce records when assessing headcount growth. Leave a field unknown when the necessary evidence is absent.

The fictional dashboard illustrates what can be measured responsibly: redundancy, attribution, recency, unresolved claims, and advertised function mix. Those measures describe the evidence process and observed source set. Their value is that a reviewer can reproduce them and see their limits.

A defensible screen ends with a bounded answer and a clear next step. It should never need a posting count to stand in for employees, revenue, business health, or acquisition readiness.

Source links

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

1. https://hiringlab.indeed.com/indeed-data-faq-2
2. https://www.bls.gov/jlt/jltdef.htm
3. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/03/how-well-do-online-job-postings-match-national-sources-in-large-english-speaking-countries_e67205a2/c17cae09-en.pdf
4. https://kb.lightcast.io/en/articles/6957661-how-does-lightcast-handle-duplicate-postings
5. https://theirstack.com/en/docs/data/job/sources
6. https://jobs.signalsapi.com/method
7. https://developers.google.com/search/docs/appearance/structured-data/job-posting
8. https://www.nber.org/papers/w32320
9. https://www.abs.gov.au/statistics/detailed-methodology-information/concepts-sources-methods/labour-statistics-concepts-sources-and-methods/2023/concepts-and-sources/job-vacancies
10. https://www.pandorlabs.com/products/job-postings
11. https://schema.org/JobPosting
12. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/bulletins/jobsandvacanciesintheuk/august2025
13. https://schema.org/totalJobOpenings
14. https://docs.greenhouse.io/job-board
15. https://www.bls.gov/news.release/jolts.tn.htm
16. https://kb.lightcast.io/en/articles/7934193-total-job-postings
17. https://www.bls.gov/web/empsit/cesfaq.htm
18. https://www.abs.gov.au/statistics/detailed-methodology-information/concepts-sources-methods/labour-statistics-concepts-sources-and-methods/2023/concepts-and-sources/jobs
19. https://www.census.gov/programs-surveys/bds/about.html
20. https://www.gleif.org/en/lei-data/access-and-use-lei-data/level-2-data-who-owns-whom
21. https://kb.lightcast.io/en/articles/7934208-unique-job-postings
22. https://arxiv.org/html/2406.06257v1

23. https://esco.ec.europa.eu/en/structure-esco-downloadable-datasets
24. https://www.onetcenter.org/content.html
25. https://www.onetcenter.org/database.html
26. https://docs.lightcast.io/lightcast-api/reference/glossary-uk-job-postings
27. https://docs.lightcast.io/lightcast-api/docs/investment-support
28. https://www.indeed.com/news/releases/indeed-turns-data-into-insights
29. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/03/how-well-do-online-job-postings-match-national-sources-in-european-countries_5e6fb2bd/e1026d81-en.pdf
30. https://www26.cedefop.europa.eu/en/publications/6217
31. https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/methodologies/vacancysurveyqmi
32. https://www.sec.gov/newsroom/press-releases/2020-192
33. https://www150.statcan.gc.ca/n1/pub/75-514-g/75-514-g2023001-eng.htm

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.