

ENTITY RESEARCH

Party and Attorney Track Record

Assemble a defensible litigation history for a party, firm, or attorney, across any Docket Navigator library.

HOW TO USE THIS DOCUMENT

Upload this PDF to the AI assistant you use with the Docket Navigator connector and ask it to create a skill from it. The pages that follow are the complete skill: the numbered sections are its instructions, and its frontmatter is listed at the end. The assistant needs the Docket Navigator MCP connector connected and authenticated before the skill will run.

Assemble a defensible litigation history for one entity: a party, a firm, or an attorney. The output is an advisory brief, not a search dump.

This workflow exists because the naive approach fails in predictable ways. A name search returns a number that looks authoritative and is usually incomplete. Remedy rows double-count. District court and agency outcomes on the same asset get flattened into one verdict. Appeal posture in the index lags reality. Each of these produces a confident, wrong answer. The steps below exist to catch them.

Library independence

Every library exposes the same shape of objects under library-scoped target names. Do not hardcode target names, join paths, or profile tab indices from memory or from another library. Establish them at the start of each run:

1. Call `docnav_plan_search_workflow` with the library name and the user's actual question. It returns candidate features, required lookups, and correct join paths.
2. Call `docnav_discover_search_options` with `library_name` and a specific `target_hint` when a target or filter is uncertain.

3. Read tab titles from `docnav_open_profile` output. Tab indices differ across libraries and profile types. Match on title, never on a remembered number.

If a lookup returns nothing, that is usually query phrasing, not absence. Rephrase toward the product's own vocabulary before concluding a concept is unsupported.

Step 1: Resolve every name variant

This is the step that determines whether the final numbers are right, so do it before anything else.

Call `docnav_lookup_search_values` on the entity criterion with the bare name. Then call it again with the distinctive fragments: the operating name, the former name, the d/b/a, the parent, the abbreviation. Entities routinely hold two or three records — one for the formal name, one for a d/b/a, one for a predecessor. Each carries different cases.

Search each variant separately and report the combined population with the split shown. A single-record count presented as complete is the most common failure in this workflow.

Note that some variants place the entity in the opposite role. A predecessor record may show the entity as a defendant in old cases. That is real history and belongs in the brief, labeled as such.

Step 2: Use the profile binder as the spine

Call `docnav_open_profile` with the resolved entity id. The profile already organizes what would otherwise take a dozen hand-built searches, with correct join paths and reliable counts.

Then call `docnav_inspect_binder_tab` on the tabs that answer the user's question. Typical tab titles and what each is good for:

TAB TITLE	ANSWERS
Cases	Full case population, forums, dates, status
Asserted / Accused [assets]	What the entity enforces or is charged with
Determinations	Merits rulings on validity, infringement, enforceability, by asset
Remedies For / Against	Monetary and injunctive awards, with appeal data
Accusations & Outcomes	Per-assertion resolution, including settlements
Claim Constructions / Terms	Construed terms and definitions
Pleadings & Orders	Motion-level history with coded results and annotations
Experience / Cases by Year	Filing trend by year and forum
Timeline, Time to Milestones	Pendency and stage timing
Firms, Affiliates	Representation and corporate relationships

Build a raw search only for a metric no tab covers.

Step 3: Keep the populations separate

Split the case list by forum before interpreting anything. District court, agency or administrative proceedings, and appellate cases are different populations with different meanings, and the entity's role often flips between them.

The asserting party is not always the plaintiff. In agency proceedings the entity is frequently the respondent to a challenge it did not initiate. Report the role per forum. A combined "cases" figure that mixes a party's own suits with petitions filed against it is not a track record.

Step 4: Do not sum the remedy column

Remedy rows are per-order, not per-award. One award commonly appears across several rows: the verdict document and the judgment document, then again on an amended judgment, plus separate rows for interest and ongoing relief.

Group remedy rows by award, then report the layers in order:

- what the factfinder awarded
- what the court entered after post-trial motions
- what interest and ongoing relief were added
- what the operative judgment currently says

Give the current operative figure and show the path to it. The headline number is what opposing counsel will cite; the operative number is what the client is exposed to. Both belong in the brief, clearly distinguished.

Also keep the object types distinct. A determination, a remedy, a claim construction, a coded motion result, and a case outcome are separate records that answer separate questions. Substituting one for another produces subtly wrong findings.

Step 5: Distinguish outcomes on the same asset

The same asset can be upheld in one forum and cancelled in another, or upheld on one ground and lost on another in the same order. When a determination table shows apparent contradictions, that is usually signal, not error.

For each asset that has been challenged more than once, report forum, date, ground, and result on one line each. Note explicitly where a challenge failed on the merits versus where it failed on a discretionary or procedural basis, because those carry opposite strategic weight. The result field alone does not tell you which happened. Open the decision.

Step 6: Read appeal posture, then verify it

Determination and remedy rows carry appeal data with appellate case numbers, dates, and results. Read it rather than assuming a judgment is final.

Then verify the current status independently. Appellate posture and case terminations change faster than any index refreshes, and this is the single field most likely to be stale on a judgment that matters. Confirm before telling a client a judgment is on appeal, final, or unreviewed. Distinguish a dismissal for lack of jurisdiction, a dismissal by agreement, and a decision on the merits — they have entirely different precedential value.

Step 7: Substantiate the decisions that carry the analysis

For any ruling the recommendation rests on, read the source rather than the coded summary. Coded event labels are indexing conventions and occasionally describe a ruling's form rather than its substance. When the source and the label disagree, follow the source and say so.

Use `docnav_resolve_product_link` with `link_kind: "filing"` and the document GUID to give the user openable links for everything cited. Do this for the full reading list, not just the top item.

Output structure

Use this shape, adapted to the question asked:

Bottom line

Two or three sentences. The finding that changes the recommendation, stated first.

Case population

Table: matter, forum, adjudicator, filed, status. Note the variant split and the combined total. State the entity's role where it varies.

What they assert

The assets, grouped by family or theory, with issue dates.

What has actually been decided

Merits rulings by forum and asset. Separate merits losses from procedural ones.

Money

The layers, in order, ending with the current operative figure and its appeal status.

Where they have lost

An explicit, enumerated list. Counsel asked for the adverse record and it is the part a favorable-looking summary buries.

Filing activity

The trend, with the inflection point identified.

Read first

Four to eight decisions in priority order, each with one line on why, each linked.

Open items

What could not be verified from the index and needs independent confirmation.

Efficiency

- Counts before rows. Tab inspection returns a count independently of the preview; use it to size the population before paging.
- Page deliberately. Pull a second page when the count exceeds the preview limit, not reflexively.
- Prefer one profile tab over three constructed searches.
- Validate incrementally when building a raw search: one criterion, confirm, then add the next. A validator rejection with a repair hint costs one call; a silently misapplied filter costs the whole analysis.
- Spend the saved calls on reading the two or three decisions that actually drive the advice.

Traps

A clean count is not a complete count.

Check variants first.

Zero results in an analytics tile can mean the activity sits under a different entity record.

Cross-check trend data against the case list before presenting it.

Derived or projected columns in analytics output

(year-to-date figures, forward estimates) are not appropriate for client-facing work unless you can state exactly how they are calculated.

An index reflects the docket at last refresh.

For anything time-sensitive — current status, appeal posture, whether a case is still live — verify independently and date the answer.

Some material facts are not in the index at all.

Asset term and expiration, corporate and supplier relationships, and settlement terms usually require outside sources. When one of these drives the recommendation, say where it came from and flag what the user should confirm.

Do not present a coded result as a holding.

"Denied" covers a merits denial and a discretionary denial. "Granted" on a motion coded one way may reflect relief the court described differently. Read the order.

Worked pattern

A user asks what an opposing plaintiff has obtained in court before advising on settlement.

Resolve the name and find two records, so the population is 31 plus 1 rather than 31. Open the profile and pull Cases, Asserted assets, Determinations, Remedies, and the trend tab. Split into district, agency, and appellate. Find 27 determinations showing one asset cancelled at the agency and a sibling upheld both there and at trial. Find 11 remedy rows that resolve to four awards: a verdict figure, a reduced judgment, interest, and ongoing relief. Read the appeal data, then check current status and discover the appeals resolved months ago, so the judgment was never reviewed. Read the post-trial opinion and find the reduction was a merits ruling rather than the remedy-type label suggests. Deliver the brief with the operative figure, the adverse rulings enumerated, and five linked decisions in reading order.

The tabs produced the structure. The verification and the source reading produced the advice.

Skill frontmatter

These fields are the skill's YAML frontmatter. Everything before this section is the skill body.

name	party-and-attorney-track-record
description	Build a complete litigation track record for a party, firm, or attorney using the Docket Navigator MCP, across any library (patent, trademark, copyright, trade secret, antitrust, securities, SEC, UPC, omnibus). Use this whenever someone asks what an opposing party, plaintiff, defendant, law firm, or attorney has actually done in litigation — their case history, win/loss record, what they have asserted, what has survived challenge, what money has been awarded, how a judgment fared on post-trial motions and appeal, or how their filing activity has changed over time. Trigger it for phrases like "what's their track record", "have they won before", "who else have they sued", "give me the full picture on this plaintiff", "how has this firm done", "research opposing counsel", or any request to assess an adversary before advising on strategy or settlement. Also use it when a single-entity question is likely to need cross-forum synthesis, because raw search results systematically undercount and mislead if assembled naively.
