

DAILY MONITORING

Docket Report Monitoring

Check a Docket Report against a watch list of parties, properties, judges, firms and attorneys, and write a ranked brief of the day's matching activity.

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.

Check a Docket Report against a watch list and produce a ranked brief of what touched it.

The watch list is the only source of what to look for. The report and the underlying records are the only sources of what happened. Never substitute your own judgment for either.

This skill is library-agnostic. Nothing in it assumes a particular Docket Navigator library, forum, section catalog, or entity vocabulary. Read those from the connector on every run.

The watch list

A spreadsheet the user maintains. It is the control surface: they should never have to open it to change it, and every run reads it fresh.

COLUMN	MEANING
Entity	The name or number as it appears in litigation. Identifiers are digits only, no punctuation.
Type	party, property, judge, firm, or attorney. Decides which part of each entry is checked. property covers whatever the library registers: patents, applications, marks, designs.
Relationship	Why the user cares, in their words. Quote this sentence verbatim in the brief to explain why an item surfaced.
Priority	high or normal. See the lift rule under Rank.
Scope	Optional forum or court abbreviation. Narrows a noisy entity to venues the user cares about. Blank means no limit.
Active	Yes or No. Skip every row that is not Yes.
Notes	Alternate captions, affiliate names, f/k/a and d/b/a forms, connector IDs, related case numbers. This column carries the variants matching depends on.

If no watch list exists, offer to build one from the entities the user names, filling Relationship from their own words.

Also read any Settings the workbook carries: which library to open, who the brief is addressed to, a global scope limit, a cap on numbered items, whether to resolve counsel, and what to write when nothing matches. Obey those values rather than hardcoding any of them.

Step 1 · Establish the library

The library determines the report, the section catalog, the forums, and what property means. Take it from Settings, or from the user, or ask. Do not guess it from the entity names.

If the library's vocabulary is unfamiliar, call `docnav_discover_search_options` or `docnav_plan_search_workflow` once to learn what it exposes. Do not invent filters, section names, forums, or proceeding types, and do not conclude a library lacks something because you have not seen it. Say a thing is unconfirmed rather than filling the gap.

Step 2 · Pull the whole report

Call `docnav_open_docket_report` with the library and no `section_names`. This returns the headlines, the publication ID, and the section catalog with a total for each section.

Then retrieve **every entry in every section listed in `available_section_names`**.

- **The section catalog is not fixed.** It varies day to day within a single library. Drive off `available_section_names` on every run, never off a list you remember or a count you saw before. A routine that hardcodes section names is broken and will not say so. When a section appears that you have not seen before, retrieve it and flag it in the brief.
- **Limit is capped at 10 per section.** Any section whose total exceeds 10 needs further calls using `start`. An eleventh entry in an eleven-entry section is invisible without pagination.
- Reconcile the retrieved count against each section's total and show the arithmetic. Do not sample.

Issue the section calls in parallel; they are independent. One overview call plus one call per section, plus one extra per additional page, is the whole retrieval budget.

To re-run a report already pulled in the same session, call the overview once and compare `publication_id` and the section totals. If both are unchanged, the publication has not been revised: re-match against what you already retrieved rather than paginating again, and say that is what you did.

What the entries actually contain

This matters more than anything else in this skill. **New-filing entries are rich. Ruling entries are nearly empty.**

New-filing entries carry the parties with their roles and IDs, the asserted properties, judges, claims, case name and number, forum, a narrative summary of the complaint or petition, and a representatives block keyed by connector firm ID. Flags on a party may mark it as a highly assertive entity; flags on a representative may mark it as local counsel.

Ruling entries carry the event with its action name, result, and the annotations holding the court's quoted reasoning, plus determinations, awards, remedies, injunctions, construed terms, and a judge field. They carry **no party list, no counsel, and often no property data at all**. The only party signal is the case name string, and that string truncates co-parties as "et

al". The events collection sometimes arrives as an object with numeric string keys rather than an array, and those keys can be non-contiguous; a gap means the publication did not carry that event, so note it and do not invent it.

Agency and tribunal entries carry the petitioner and rights-holder names, the property number, and any institution or outcome data, without a party array.

Attorney names are always omitted from report payloads, in every section.

Step 3 · First-pass match

Check every retrieved entry against every active row. Apply the row's Scope and any global scope limit.

- **party** — check every party field the entry carries, in every role. Match affiliate captions and variants from Notes. Always state the role: at a tribunal the user's client may be the petitioner rather than the respondent.
- **property** — exact identifier only, digits compared to digits. Check asserted properties, determinations, construed terms, and any property number field. Family members, continuations and related registrations are never matched.
- **judge** — check the judges collection and the judge field, and apply Scope.
- **firm** — representatives keys are connector firm IDs, so firms match by ID with nothing expanded. Notes carry IDs for known firms.
- **attorney** — impossible at this stage. Defer to Step 4.

This pass is triage that finds candidates. It is not sufficient on its own, and a non-match here means very little.

Step 4 · Open the underlying records

This is where most real matches are found. Use `docnav_open_profile` with `profile_type`: "case" and the case ID. Open the record for:

1. **Every ruling entry, matched or not.** Ruling payloads have no parties, no counsel and often no properties, so a Step 3 non-match on a ruling carries no information. Skipping the unmatched ones is how a client gets missed.

2. **Every entry that matched in Step 3.**

3. **Every entry carrying a representatives block**, when a firm or attorney row is active.

A new-filing entry whose payload already supplied parties, properties and a judge needs no record opened unless it matched or unless counsel rows are live. That is where the efficiency is.

Parse the header field. It is docket header markup carrying the full party list, every attorney and firm with role and termination date, and the judge assignment line. Also read the structured property numbers, judges, accused and asserting parties, and any termination date. Re-run **all five row types** against everything you find.

What the records catch that the report does not

- **A ruling caption names one side, and "et al" hides the rest.** A co-defendant client concealed behind "et al" is the single most common miss, and the payload will never name them.
- **Ruling entries carry no counsel at all.** A watch-listed firm appearing as opposing counsel is invisible until the header is read.
- **Ruling entries often carry no property data.** Re-run property rows against the record's property numbers.
- **New-filing entries can show an empty judges collection while the docket has an assignment.** Check the record before reporting a case as unassigned, and distinguish genuinely unassigned from missing.
- **Tribunal and agency case profiles return a null header.** There is no docket header to parse, so firm and attorney rows can never resolve on those proceedings from this source. Their structured party arrays are populated, so parties remain checkable. State that limitation rather than implying counsel was checked.
- **Grouped or schedule-based defendant blocks are not itemised.** A watch-listed party could sit inside one unverifiably. Say so rather than reporting a clean negative.

Where a fact came from a record rather than the report, say so in the brief.

If Settings disables the counsel lookup, that suppresses **only** firm and attorney resolution. Still open records, and still run party, property and judge matching against them. Party identity on a ruling is not a counsel question.

When a ruling entry looks thin

Trigger this for any ruling with no events, a null judge, an award or determination with no annotation explaining it, or any disposition whose reasoning is absent. Also for any Priority 1 ruling, since those carry the most consequence.

Run a bounded event search with `docnav_validate_and_preview_search` against the library's event target, filtered to the case ID and a date window of roughly ninety days ending at the report date. Rows return the action name, result, filing date, document ID and document GUID. The returned count is coded document-type-events rather than distinct orders, so report the distinct-document figure instead.

To read a document, pass the bare UUID to `docnav_read_filing_text`, stripping any API path prefix from the GUID. For a long order, read the opening pages for posture and the closing pages for the holding rather than the whole thing.

This answers what the publication leaves open: what produced the ruling, who filed it, whether it was opposed, who signed it, and what judgment it sits on. A publication can carry one document out of a multi-document event and omit the order that explains it, and a thin costs or fee entry can sit on a very large judgment that never appeared in any report.

Confirming a firm

Use `docnav_open_profile` with `profile_type: "firm"` to confirm what a firm ID belongs to. Treat its spellings list as **confirmation only, never as a matching source**. Those lists carry former names that collide with unrelated firms, and matching on them produces false positives. Match on ID, or on a full name read from a docket header.

Watch for near-miss names generally. An attorney sharing a surname with a watch-listed judge is not a match; note the near miss rather than silently dropping it.

Step 5 · Rank

- **Priority 1** — a direct client or active-matter match. A client is on the caption, or a ruling issued in a matter the user is in.
- **Priority 2** — a direct tracked party, property, firm or attorney match outside the user's own matters.
- **Priority 3** — contextual, such as a tracked judge ruling in a case the user is not in.

A client on the caption satisfies Priority 1 whether or not the user is counsel. When a Priority 1 item is one they have no role in, say so plainly in the item so they are not left inferring an engagement that does not exist.

The high lift. A high row lifts its item one level: Priority 2 becomes Priority 1, Priority 3 becomes Priority 2. On something already at Priority 1 it does nothing. When several rows fire on one item, the item takes the highest resulting priority and the brief lists every row that fired. Note in the coverage section when a lift alone put a non-client matter into Priority 1, so the user can retune the flag if Priority 1 stops meaning client work.

Firm-only and attorney-only matches, where no tracked party, property or judge is on the caption, do not get numbered items. Collect them in a short list at the foot of the brief: firm or attorney, case name, forum, number, date, parties, and one line on what was filed. This keeps numbered items tied to the user's own entities.

The item cap. When matched events exceed the cap in Settings, group by entity: several proceedings involving one tracked entity become a single numbered item with the events as bullets. Never drop a Priority 1 item to fit; if grouping is not enough, keep Priority 1 and 2 and compress Priority 3 to one line stating how many were compressed. Report each real-world event once, and report a new event in a matter already seen.

Within each level, order by consequence. Anything adverse in an active matter leads.

Step 6 · Write the brief

A markdown file in the user's folder, named for the **report date** rather than the run date. One file per report date; a re-run replaces it.

- Open with the matched-item count, the single most serious thing, and a review counter giving entries and sections reviewed against matches found. Take the section count from what the run actually retrieved. Keep the counter neutral: entries that did not match are non-matches, not "dismissed" or "unrelated".
- Group under Priority 1, 2 and 3 headings. If a level is empty, say so.
- Per item: case name, forum, number, judge, date; the rows that fired with their Relationship sentences quoted; then what happened. Quote the court's own language where the annotation carries the reasoning. Close with what it means for the client or the matter.
- Give every item a link back to Docket Navigator, using the profile and report URLs the connector returns rather than any URL you construct. **Never build a link to an external docket system.** Identify motions and briefs by docket number only.
- Size each item to its consequence, judged by what Step 4 found rather than by what the publication printed. A costs entry sitting on a large adverse judgment is not a routine costs entry. Do not pad items to match each other.
- Where a chain was reconstructed, give it in order ending at the ruling, each row carrying docket number, date and filer. Attribute the parties' positions to the order's own recitation. **Never attribute an argument to a brief you did not read,** and never claim to have read one. Where no chain exists, say so rather than assembling a plausible one.
- Close with coverage notes: rows suppressed by Active = No and what each suppression actually cost that day, verified case by case rather than asserted; negative confirmations naming the rows that did not match; any high lift promotions; the exact-identifier limitation where a ruling turns on an unlisted family member; the null-header limitation wherever tribunal or agency entries appeared; any section outside the usual catalog; anywhere the publication omitted a document the coded events carried; and unresolved discrepancies between the report and a record.
- If nothing matched, write the zero-match line plus the review counter. That line is the evidence the check ran. Do not pad it with near-misses.

Present the file, then summarise in chat only the two or three things needing attention, leading with anything adverse in an active matter. Do not restate the brief.

Step 7 · Verify before presenting

1. Retrieved entries equal the sum of section totals, across every section in the catalog.
2. Every active row is accounted for as a match or a deliberate non-match. Walk the list.
3. Every ruling had its record opened, including the ones that did not match on the payload.
4. Every thin ruling and every Priority 1 ruling had its events pulled, or the brief says why not.
5. Every number, name, holding and amount traces to a specific tool response.
6. Procedural posture is exact. A recommendation is not an order, an institution decision is not a final decision, a stay is not a dismissal, a clerk's action is not a judge's ruling, and nothing is invalid or cancelled unless a determination or annotation says so.
7. Any structured field naming a single payer, recipient, petitioner or defendant is checked against the underlying document before the brief repeats it. These fields have named one entity where the judgment ran against two.
8. Suppressed rows: confirm what each would have surfaced so the brief can state the cost.
9. Where the report and a record disagree, report both and mark it unresolved rather than silently choosing.

Editing the watch list

Users change the list conversationally. Treat it as a write to their file.

Confirm before writing, even when the instruction is explicit. Set Active to No rather than deleting a row: deletion loses the Relationship text, which is the record of why the entity was ever tracked and the sentence the brief quotes. Record the reason the user gave in Notes. Report which cell in which row changed.

Afterwards, reopen the workbook and confirm that every other row is untouched, that all sheets survive, and that any dropdown validation still works. Verify by reading the file, not by trusting your own report of the edit.

Then re-read the sheet before re-running. Never assume your own edit landed.

When a re-run follows an edit, state how the item count changed and which items moved. Where a paused row still matched on surviving items, say so: that tells the user the row was confirming rather than discovering. Report what changed without adopting a view about the entity that came off the list.

Scheduling

When the user wants this to run on a cadence, create one recurring scheduled task bound to their folder, confirm before creating it, and report the next run time. Keep the task prompt a thin wrapper that invokes this skill, so the skill stays the single source of truth and later fixes reach the run that matters. The prompt must say "today's report" and must never pin a date; a task carrying a fixed date runs forever against one stale publication and reports success.

Tell the user that connector access needs to be approved before the first unattended run, and note that a small offset between the scheduled time and the displayed run time is normal load spreading rather than an error.

Accuracy rules

Everything factual comes from a tool response. Do not infer counsel or parties from a narrative summary when a record is available. Do not describe anyone as a client, prospect or adversary beyond what the Relationship column says. If something is unconfirmed in the metadata, say it is unconfirmed. Report what the sheet asks for even when unwelcome, and say plainly when a ruling goes against the user's client.

Forum is a scope setting, never a trigger; tracking a busy court on its own returns that court's entire daily volume. Legal issues, doctrines, motion types and outcomes are never triggers either. They are detail on items that already matched an entity.

Skill frontmatter

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

name	docket-report-monitoring
description	Check a Docket Navigator Docket Report against a watch list of parties, properties, judges, firms and attorneys, and write a ranked brief of the day's matching activity. Works with any Docket Navigator library. Use when the user asks to run their docket check or watch list, produce a morning docket brief, monitor a report for their clients or adversaries, or asks what in today's Docket Report touches their matters.
