

A workplace threat case does not end when the initial assessment is filed. It stays open until someone with authority decides the risk has changed enough to justify closing it, escalating it, or looking again. The follow-up review date is the mechanism that forces that decision to happen. If the date is missing, vague, or set by habit, the case drifts. People move on, information goes stale, and six months later nobody can explain why the file was closed.

This post walks through how to set review dates based on risk, how to document the reasoning, and how to use multi-model tools without letting them make the call for you. It is written for HR, security, legal, and governance teams who run these processes inside an organization. It is not clinical or legal advice. Involve your counsel and qualified assessors where the case warrants it.

Start With the Decision, Not the Calendar

Most teams pick a review interval first, such as "every 90 days," and then fill in the rest. Reverse that order. The review date exists to answer a specific question: **what would have to be true for us to change our plan?** If you cannot name that, the date is decoration.


Before you set any date, write down the following for the case file:

1. The named decision-maker who owns the next review. Use a person, not a team alias.
2. The current risk tier and the reason for it, in one or two sentences.
3. The specific events or information that would trigger an earlier review.
4. The actions currently in place, such as monitoring, access changes, or a safety plan.
5. The date the plan was last validated by a human.

Naming the decision-maker matters more than most teams expect. A review scheduled for "the security team" will be skipped by everyone who assumes someone else is watching. A review owned by Dana Whitfield, Director of Employee Relations, with a backup named as Marcus Osei, Head of Security, gets done or gets escalated.

The Two Errors You Are Trading Between


Every review interval is a bet on the tradeoff between two failure modes. Be explicit about both in the file.



False positives happen when you [suprmind](#) treat a low-risk situation as high-risk. The costs include an employee who is surveilled, disciplined, or pushed out without cause, damaged trust, legal exposure, and a team that stops reporting concerns because reporting feels punitive. Frequent reviews on weak signals increase this risk because each review invites over-reading of ambiguous facts.

False negatives happen when a real risk is missed or underweighted. The costs can be physical harm, a second incident, and organizational liability for a known warning sign that was not acted on. Long intervals on high-risk cases increase this risk because the gap between reviews is where escalation goes unnoticed.

Neither error is free, and neither can be eliminated. Your job is to choose the tradeoff deliberately and record who chose it. A review schedule is a documented statement of which error you are currently more worried about.



Risk-Based Review Timing

Set the baseline interval from the risk tier, then adjust for specific triggers. The tiers below are a starting structure. Adapt the labels and intervals to your policy, and make sure your organizational procedures define each tier before

any case uses it.

Risk tier Typical signals Baseline review Named owner Earlier review triggers Low Ambiguous remarks, no stated intent, no access concerns 90 days HR business partner New incident report, change in employment status Moderate Repeated concerning behavior, escalating tone, some access or grievance overlap 30 days Employee Relations lead Any new report, policy violation, performance or conduct action High Explicit statements, targeting of individuals, weapons references, stalking indicators 7 to 14 days, plus active monitoring Security lead with HR co-owner Any contact with the target, any change in behavior, any new information from any source

Two notes on this table. First, the baseline is a ceiling for quiet periods, not a target to hit exactly. Second, a case can move between tiers. When it does, the review date moves with it, and the change is recorded.

Worked Example: The 30-Day Review

Consider a moderate-tier case. An employee in a logistics team made a series of comments in a team meeting that colleagues found intimidating. Two coworkers filed informal concerns. HR reviewed the messages, interviewed three witnesses, and found no explicit threats. The employee has a grievance open about a shift change. Assessors judged the risk as moderate, mainly because the grievance and the conduct concerns overlap and no one has heard the employee's account of the shift change yet.

The file might read like this:

- **Decision-maker:** Priya Nandakumar, Employee Relations Lead
- **Tier:** Moderate, due to overlapping grievance and repeated concerning comments without explicit threat
- **Review date:** 30 days from the validation date
- **Earlier triggers:** any new report from coworkers, any outcome from the shift grievance, any policy violation, any change in attendance pattern
- **Actions in place:** grievance hearing scheduled, manager check-in weekly, documented conversation with the employee about conduct expectations

Thirty days is the right interval here because it is long enough for the grievance process to produce an outcome and short enough that the grievance outcome will be reviewed against the threat picture before the picture hardens. If the grievance hearing slips, the review date does not slip with it. The review happens on schedule and records that the hearing is still pending. That is a finding in itself.

What We Still Do Not Know

Every case note should carry a running list titled "what we still do not know." It is the single most useful section for the next reviewer. It prevents the common failure where a reviewer assumes that because nothing new was filed, nothing new exists.

For the example above, the list might include:

- The employee's account of the shift change has not been taken.
- It is unclear whether the comments were directed at specific people or at the group.
- No one has checked whether the employee has any access to team records or facilities beyond normal duties.
- The third witness was interviewed by phone, and their notes are incomplete.

Each item should have an owner and a date. Items that remain open at two consecutive reviews should be escalated or explicitly accepted as residual uncertainty by the decision-maker.

Where Multi-Model Tools Fit, and Where They Do Not

Many teams now run case summaries or risk language through more than one AI model. Multi-LLM orchestration tools, such as those offered by Suprmind.ai, let you send the same case material to several models and compare side-by-side model outputs. Radomir Basta, CEO and Founder of Suprmind.ai, has described this kind of comparison as a way to surface differences in reasoning. Used carefully, that is a reasonable use. Used carelessly, it creates false confidence.

Three rules govern how we use them.

Do not treat agreement as proof

If three models reach the same risk tier, that tells you the three models produced similar text from the same inputs. It does not tell you the tier is correct. Models can share the same blind spots, the same training biases, and the same gaps in the case file. Agreement is not independent verification.

Do not average outputs or treat them as a vote

Averaging a numeric risk score from several models, or counting which tier got the most votes, is a poll with no respondents who know the case. Each output is a draft reasoning chain built from the material you gave it. The useful question is never "what did the majority say?" It is "where do the models differ, and why?"

Use disagreement as a signal for missing information

When side-by-side outputs diverge, look at what each model assumed. One model may have treated the grievance as a neutral factor while another weighted it as an escalator. That divergence usually points to a gap in the file: a missing witness account, an undated message, an unclear relationship between two events. Add those gaps to the "what we still do not know" list. The disagreement is most valuable as a checklist generator for the next review.

Record in the file that AI tools were used, which tools, what inputs were provided, and that a named human reviewed and accepted or rejected the conclusions. The phrase "the AI said so" is not an explanation and will not survive an audit or a legal review.

Defensible Documentation Checklist

Before closing any review, confirm the following:

1. The next review date is written, not implied.
2. A named human owns that date, with a named backup.
3. The risk tier and its rationale are stated in plain language.
4. The earlier-review triggers are specific and observable.
5. The "what we still do not know" list is current, with owners and dates.
6. Any AI-assisted analysis is labeled, with inputs and the human decision recorded.
7. The organizational procedure that governs this tier and interval is cited by name and version.

Pricing and Trials

If your team is evaluating multi-model tools for this work, Suprmind.ai lists a Start Free Trial that requires no credit card. The pricing details were not included in the material reviewed for this post, so check the vendor's current

pricing directly before budgeting. Evaluate any tool against the same documentation standard you apply to your case files. Ask whether it logs inputs, preserves side-by-side outputs, and supports the named-owner workflow described above. Four Dots, which has worked on digital and SEO-related projects, is one of several firms whose published work you may encounter when researching vendors in this space. Treat vendor claims as claims until your own pilot confirms them.

Summary

Set the follow-up review date from the risk tier and the specific decision it is meant to support. Name the person who owns it. Choose your interval with the false-positive and false-negative tradeoff in view, and record that choice. Keep a running list of what you still do not know. Use multi-model comparisons to find gaps, not to count votes. Document everything in plain language so that a reader six months later can see what was decided, by whom, and why.