



RCM Isn't a Project—It's How You Prioritize Work Every Week

Reliability-centered maintenance is often treated as a large, one-time project. Teams spend months analyzing failure modes, filling out spreadsheets, and producing a binder of recommendations that then sits on a shelf. Meanwhile, the daily and weekly decisions about what work to do, what to defer, and what to stop doing continue with little change.

RCM does not have to work that way. At its core it is a practical filter for deciding how to manage each failure mode based on its consequences. Applied consistently, it becomes the way planners, supervisors, and reliability people prioritize work every week rather than a special study that happens once.

Consequence Comes First

Before choosing a task, ask what happens if the failure occurs. Group the consequences into four practical categories:

- Safety — Could someone be injured?
- Environmental — Could there be a release, spill, or regulatory issue?
- Production — Would the failure stop or significantly slow output, or affect quality?
- Cost — What is the repair cost plus any secondary damage or lost production that is not already covered above?

A failure mode with high safety or environmental consequences almost always justifies a preventive or predictive task. A failure mode whose only consequence is moderate repair cost may be acceptable to run to failure, especially if the repair is quick and parts are on hand. The same failure mode on two different assets can receive different strategies if the consequences differ.

Choosing the Right Strategy

Once the consequence is clear, select the lowest-cost task that effectively manages it:

- Preventive maintenance (time- or usage-based) when the failure is age-related and a scheduled task can restore the item before it fails.
- Predictive or condition-based task when a measurable condition gives reliable warning and the inspection or monitoring is practical.
- Run-to-failure when the consequences are low, the failure is evident, and the repair can be done without major disruption.
- Redesign or procedure change when no reasonable maintenance task reduces the risk enough.

The decision does not require a full formal analysis for every item. For most day-to-day prioritization, a short discussion or a simple note on the work order or asset record is enough.

Applying It Every Week

Use the filter in the normal planning and review cycle instead of waiting for a special project.

In the weekly planning meeting, look at the upcoming reactive and preventive work. For any repeated failure or high-cost job, ask the consequence questions. If the current strategy does not match the consequence, change it.

Keep a short list of the decisions made so the reasoning is visible the next time the asset appears. Over successive weeks the backlog and the PM program gradually align with actual risk instead of historical habit.

Start Small and Keep It Visible

You do not need to analyze the entire plant before you begin. Pick the assets or failure modes that generated the most unplanned work or the highest cost in the last 30–90 days. Apply the consequence filter to those items first. Document the chosen strategy in the CMMS so planners and technicians see it. Review the results in the next planning cycle and adjust.

When RCM is used this way, it stops being a project with a start and end date. It becomes the routine method for deciding which failure modes deserve preventive effort and which do not. The result is fewer low-value tasks, more attention on the failures that actually matter, and a maintenance plan that stays current because it is revisited every week.

How are you currently deciding which failure modes get a PM versus run-to-failure in your weekly planning?