

Mitigation Prep Standard: How to Build One That Works

This is the second article in the Mitigation-to-Reconstruction Intelligence cluster under [The Restoration Operator's Playbook](#). It builds on [the handoff piece](#) — read that first if you haven't.

The standard is the moat

If the mitigation-to-reconstruction handoff is the most expensive moment in restoration, the documented mitigation prep standard is the operational artifact that converts that expense into an advantage. It is also the artifact that almost no one in the industry actually has.

Operators talk about prep standards all the time. They mean different things by the phrase. Some mean a set of unwritten norms that the senior crew carries in its head. Some mean a few pages in an employee handbook that nobody references after the first day of orientation. Some mean a software workflow that captures dryout readings and calls itself a standard. None of those are the thing.

The thing is a written, version-controlled, operationally specific document that tells a mitigation tech how to make the cut, demo, removal, and documentation decisions that have downstream reconstruction consequences. It is the single most important operational document a restoration company will ever produce, and the companies that have built one know it.

This article is a description of what such a standard actually contains, how it gets written, and why most attempts to build one fail.

What a real prep standard contains

A working prep standard is not a manual. It is a decision aid for the moments when a mitigation tech is standing in a structure with a utility knife in their hand and a sixty-second window to make a choice that the rebuild team will live with for the next ninety days. The standard has to be specific enough to produce a different decision than the tech's instinct would, in the cases where the tech's instinct is wrong.

The categories of decisions it has to address fall into a predictable pattern across most water and fire losses.

The first category is cut decisions on drywall. How high to cut. Whether to cut along a stud line or use a flood cut. How to handle the meeting points between affected and unaffected areas in a way that produces a clean rebuild seam. How to handle ceilings where the cut decision interacts with insulation and texture matching. The standard names the default choice for each of these, the conditions under which the default changes, and the conditions under which the tech is expected to call a supervisor before cutting.

The second category is removal decisions on baseboards, trim, casing, and crown molding. Whether to remove and reuse, remove and discard, or leave in place and treat. The default choice is rarely the same across all conditions — paint-grade and stain-grade trim warrant different defaults, modern composite trim warrants a third, and historical or custom-milled trim warrants a fourth. The standard documents which is which and how to identify each in the first ten minutes on site.

The third category is flooring. Where the cut line goes, how to handle transitions to unaffected areas, when to remove pad versus pad and carpet, when to remove tile versus dry in place, how to handle engineered hardwood versus solid, how to handle LVP and the specific question of whether to lift to a natural transition. This is the category where the rebuild team is most often blindsided by mitigation decisions, because flooring rebuild aesthetics are entirely a function of where the mitigation crew chose to stop cutting.

The fourth category is cabinetry, vanities, and built-ins. When to remove the kicks. When to pull cabinets entirely. When to drill weep holes. When to dry in place with cavity drying. The standard has to acknowledge that these decisions are partly a function of the cabinet construction, partly a function of how the rebuild team prefers to receive the job, and partly a function of carrier expectations. The default choices and the override conditions need to be specified.

The fifth category is documentation: photo angles, lighting conditions, what to capture before any work begins, what to capture during demo, what to capture after demo, how to label, how to organize for both the carrier file and the rebuild estimator. This is the category most undervalued by operators who have never been the rebuild estimator opening the file two days later. Documentation discipline that is built around the rebuild estimator's needs prevents the largest single source of wasted estimator hours in the industry.

The sixth category is communication: when the mitigation supervisor calls the rebuild team, when the rebuild team is brought to site, when the homeowner is told what to expect about the rebuild, who owns each conversation. Communication failures account for a surprising fraction of the friction the rebuild team encounters, and most of those failures are fixable with a written protocol about who talks to whom when.

How a real prep standard gets written

The standard cannot be written by a single person sitting in an office. It also cannot be written by a committee. The companies that have produced working standards have followed a specific pattern.

The work begins with one operator who has done both sides of the job — mitigation and reconstruction — and who has the credibility internally to make decisions stick. That operator is the author. Not a committee chair. The author. They are responsible for the document being good and for it being adopted.

The author starts not with their own knowledge but with the recent failure log. The last ninety days of completed jobs, walked one by one with the reconstruction estimator and the mitigation supervisor. For each job, the question is the same: where did the rebuild team have to do extra work, eat margin, or take a homeowner concession because of a mitigation decision? Each instance gets logged, categorized, and converted into a decision rule that, if it had been in place at the time, would have prevented the problem.

The first draft of the standard emerges from this exercise. It is not comprehensive. It is not elegant. It addresses the specific failure modes the company has actually experienced. That focus is a feature, not a bug. A standard that tries to cover every conceivable scenario gets ignored. A standard that addresses the twenty things that go wrong most often gets used.

The first draft then gets pressure-tested in two ways. The mitigation crew leads read it and challenge anything that seems impractical, slow, or based on a misunderstanding of how the work actually happens in the field. The rebuild estimators read it and flag anything that does not actually solve the rebuild problem they were complaining about. Both groups have to feel ownership before the standard ships.

Then it ships. Not as a binder. As a short, scannable document — usually ten to twenty pages — that lives in the company's operational system, is referenced in every job kickoff, and is the basis for the company's mitigation training program.

And then, critically, it gets revised every quarter. The companies that have done this for several years describe their current standard as “version eleven” or “the November rev.” It is a living document. The day it stops being revised is the day it starts being ignored.

Why most attempts to build one fail

Most companies that try to build a prep standard fail. The failure modes are predictable.

The first failure mode is committee authorship. A standard written by consensus reads like a treaty. It hedges every decision, includes too many exceptions, and produces no behavior change. The author has to be one accountable person.

The second failure mode is starting from theory instead of failure. Standards written from first principles or from industry best practices end up being too generic to change anything in the field. The standard has to come out of the company's actual recent failures, because those are the failures the field crew will recognize and accept guidance on.

The third failure mode is over-comprehensiveness. A two-hundred-page standard does not get read. A standard that addresses the twenty most common decision points and is honest about not addressing the rest is the one that gets used. Coverage is not the goal. Behavior change on the highest-value decisions is the goal.

The fourth failure mode is publishing without training. A document that is sent out with a memo gets ignored. A document that is the basis for a half-day field training, with the senior author walking the crew through each decision and the reasoning behind it, gets adopted. The training is part of the standard, not a follow-up to it.

The fifth failure mode is no revision cadence. Standards that ship and then sit on the server for two years stop matching the current state of the work. The crew learns to disregard them. A quarterly revision cycle, even if most quarters only produce small updates, keeps the document credible.

The sixth failure mode is treating the standard as the property of the operations function alone. A standard that the mitigation crew owns but that the rebuild team does not actively use as a quality scorecard is half a standard. The rebuild team has to be empowered to flag deviations, and the flags have to feed back into the next revision. Without that loop, the standard ossifies.

What the standard does to the company

The companies that have built and maintained a real prep standard for several years tend to describe similar effects. None of the effects are about the standard itself. They are about what the standard makes possible.

The first effect is on training. A new mitigation tech can be brought from green to credibly autonomous in a fraction of the time a similar tech would take in a company without a standard. The standard is the curriculum. The senior tech who would have been burned mentoring one apprentice at a time can mentor a whole class against the standard, with much higher consistency in the output.

The second effect is on rebuild margin. The rebuild estimators stop encountering the surprises that used to eat their hours. Estimates get written faster, get approved faster, and produce fewer scope arguments. The margin recapture from this effect alone usually pays for the standard work many times over within the first year.

The third effect is on customer experience. The handoff feels different to the homeowner. The mitigation crew leaves a job that the rebuild team can pick up cleanly, which means the rebuild starts faster, runs cleaner, and finishes with a homeowner who feels the company knew what it was doing the whole way through. Five-star reviews go up. Complaints go down.

The fourth effect is on the relationship with carriers and TPAs. The pattern of clean files, clean scope discussions, and rare disputes gets noticed. Program placement improves. Referral flow improves. The carrier-side reputation compounds in a way that takes years to build but is durable once built.

The fifth effect is on the company's ability to absorb new technology. A documented standard is the substrate that makes AI-assisted operations possible. Software that is asked to apply judgment to new situations performs as well as the documented judgment it has access to. Companies with a real standard can plug new tools in and get force multiplication. Companies without a standard buy tools and watch them fail to deliver, because the tools have nothing to ground their decisions in.

Where to start if you don't have one

If you run a restoration company and you do not have a prep standard, the work to produce one is genuinely hard, but the starting point is not. Pick the operator on your team who has done both mitigation and reconstruction and who has the credibility to make decisions stick. Have them block one full afternoon with the rebuild lead and the mitigation supervisor. Walk the last ten completed jobs file by file, asking the failure question described above and in [the handoff piece](#).

That afternoon will produce a list of fifteen to twenty-five recurring failure modes. Each of those failure modes is a decision rule waiting to be written. The first draft of the standard is just those rules, written down, in the voice of the author, with the conditions and the override criteria specified.

That first draft is not the finished product. But it is the artifact that, more than any other single thing the company will produce in the next twelve months, determines whether the company is on the operating-system side of the industry split described in [the pillar piece](#) — or the side that wakes up in 2028 wondering what happened.

The standard is the moat. The companies that build it know it. The companies that don't are about to find out.

Next in this cluster: photo and documentation discipline built around what the rebuild estimator actually needs to see. After that: the feedback loop that turns rebuild discoveries into the next revision of the standard, and the shared metrics that hold both teams accountable to the same scoreboard.