

Mitigation Reconstruction Handoff: The Shared Scoreboard

This is the fifth and final article in the Mitigation-to-Reconstruction Intelligence cluster under [The Restoration Operator's Playbook](#). It builds on [the handoff piece](#), the [prep standard piece](#), the [photo discipline piece](#), and the [feedback loop piece](#).

Two functions cannot share a job if they do not share a number

The hardest problem in the mitigation-to-reconstruction handoff is not technical. It is not procedural. It is not even cultural in the broad sense. It is a measurement problem.

In most restoration companies, the mitigation function and the reconstruction function are measured on different numbers. Mitigation is measured on dryout time, equipment utilization, response speed, maybe a per-job revenue or margin number specific to the mitigation portion of the work. Reconstruction is measured on cycle time, gross margin per job, scope accuracy, customer satisfaction at the close-out. Each function tracks its number, manages to its number, and gets rewarded based on its number. Each function is, in a literal accounting sense, optimizing for a different outcome.

The handoff lives in the gap between those two numbers. There is no metric that captures whether the handoff was good or bad. There is no scoreboard that holds either function accountable for the other's experience. The handoff is, by structural design, no one's number.

The single highest-leverage operational change a restoration company can make to fix the handoff problem is to put both functions on the same scoreboard for at least one number that captures the joint outcome. Not instead of their function-specific numbers — in addition to them. The shared number is what makes the prep standard, the photo discipline, and the feedback loop work in concert. Without a shared number, all three of those artifacts can exist on paper and still produce no behavior change.

What the shared number has to be

For a shared metric to work, it has to satisfy three criteria.

It has to be a number that both functions genuinely influence. A metric that is mostly driven by mitigation but slightly affected by reconstruction will be experienced by the reconstruction team as

unfair, and vice versa. The number has to be one where both teams can point to specific decisions they make that affect it.

It has to be measurable at the job level, not the function level. Function-level numbers create function-level optimization. Job-level numbers force the two functions to think about the joint outcome on each individual file. Aggregations across jobs are useful for trend reporting, but the number has to live first at the job.

It has to be visible quickly enough to drive behavior. A metric that takes ninety days to settle is too slow to influence the next decision the mitigation tech makes. The number has to close out within a window that lets both teams see the result of their handoff and adjust.

The number that satisfies all three criteria in most restoration companies is total job margin, measured at the job level, with both teams accountable to it.

Why total job margin is the right number

Total job margin captures everything that matters about the handoff. A mitigation crew that demos too aggressively raises the rebuild scope and depresses total job margin even if the mitigation portion looks healthy. A mitigation crew that documents poorly creates rebuild rework that depresses total job margin even if the mitigation portion was efficient. A mitigation crew that prepares the job well for the rebuild produces a job where both portions perform, and total job margin is high.

Conversely, a rebuild team that consistently writes scope that fits the conditions the mitigation crew left will produce healthy total job margins on jobs where the mitigation work was good and surface the handoff problems clearly on jobs where it was not. The rebuild team is also incentivized to communicate clearly with the mitigation team about what kinds of prep work consistently lead to healthy rebuilds, because better prep raises the number they are accountable to.

The mitigation team, in turn, becomes interested in what happens after they leave the job. A mitigation supervisor who sees that their jobs consistently produce lower total margins than peers' will start asking why. A mitigation supervisor whose jobs consistently produce higher total margins will be asked to teach the rest of the team. The conversation about the handoff stops being political and starts being operational.

Total job margin also has the practical advantage of being a number every restoration company already calculates. The work to put it on a shared scoreboard is mostly the work of presenting it differently — at the job level, visible to both functions, attached to the leadership review of both functions.

Secondary metrics worth sharing

Total job margin is the primary shared metric. Several secondary metrics, used in addition to the primary, sharpen the picture and make the joint accountability more actionable.

Total job cycle time — from first notice of loss to keys-back-to-homeowner — is the most useful secondary metric. It captures whether the handoff added unnecessary days to the timeline. Mitigation crews that hand off cleanly contribute to shorter cycles. Rebuild teams that pick up cleanly do the same. Both teams seeing the cycle time at the job level creates pressure to find the days that are being lost in the handoff.

Customer satisfaction at the close-out, captured through whatever survey or review mechanism the company uses, is a useful third metric. Customer satisfaction is more sensitive to the rebuild experience than the mitigation experience, but it is influenced by both, and putting it on a shared scoreboard prevents the mitigation team from optimizing purely for their own customer interaction at the expense of the longer arc of the homeowner's experience.

Scope change rate during the rebuild — how often the rebuild team has to write change orders or get scope adjustments approved — is a fourth useful metric. A high scope change rate often traces back to incomplete handoff documentation, undiscovered conditions that should have been flagged at mitigation, or decisions that should have been made differently at the front of the job. Tracking it as a shared number drives both teams to invest in the documentation and prep work that prevents it.

None of these secondary metrics replaces total job margin as the primary. They support it. They give the leadership conversation specificity when the primary number drifts in a direction that needs investigation.

What changes when the scoreboard becomes shared

The companies that have implemented shared scoreboards across the mitigation and reconstruction functions report a similar set of changes.

The first change is in conversation. The mitigation supervisor and the rebuild lead start talking to each other differently. The conversations stop being about whose fault something was and start being about how to make the joint number better. This shift is small in any single conversation and large over hundreds of conversations across a year.

The second change is in decision-making. Mitigation crews start making cut, demo, and documentation decisions with more attention to downstream consequences, because they know the consequences will

show up on a number they are accountable to. Rebuild teams start engaging earlier on jobs, sometimes visiting site during mitigation on complex losses, because the early engagement protects the joint number.

The third change is in training and hiring. The standards that govern the work get communicated as joint standards rather than function-specific standards. New hires on either side learn that they are part of a joint operation, not a siloed function. Senior operators on both sides become natural cross-trainers, because the joint number rewards cross-functional fluency.

The fourth change is in technology investment. Software and tooling decisions start being evaluated against their effect on the joint number rather than the local efficiency of one function. This usually leads to better tooling decisions, because the joint outcome is what the company actually cares about.

The fifth change is in leadership focus. Owner and senior leader attention starts following the joint number, which puts the right kind of pressure on the right kind of operational improvements. Function-specific dashboards still exist, but the joint dashboard becomes the one that drives the operating cadence.

Why most companies do not do this

The barriers are not technical. The numbers exist. The systems can produce them. The barriers are political and operational.

The political barrier is that function leaders have built their careers around function-specific metrics. Asking them to share accountability with another function feels like a dilution of their authority and a complication of their performance evaluation. The owner has to be the one who makes the call, and the call has to be made deliberately, with explicit acknowledgment that the function-specific metrics still matter and that the shared metric is additional, not a replacement.

The operational barrier is that most operations software is configured to report function-specific numbers and not configured to surface job-level joint numbers in a useful way. Producing a clean joint scoreboard usually requires either a custom report, a workaround in the existing software, or a small investment in a reporting layer that pulls from the operations system and presents the data the way the joint conversation needs to see it. The work is not large, but it has to be commissioned, and in most companies no one has commissioned it because the conversation about the joint metric has not yet happened.

The cultural barrier, which is the deepest, is that some companies have developed cross-functional dynamics over years that would be uncomfortable to surface. A shared scoreboard makes visible patterns that have been invisible. Some of those patterns will be flattering to one function and unflattering to

another. The leadership has to be ready to handle that surfacing constructively, or the scoreboard will become a weapon and the experiment will fail.

How to start

If you run a restoration company and you do not have a shared scoreboard, the path to building one is short.

Calculate total job margin at the job level for the last six months. Most operations systems can produce this with modest effort. Surface it to both function leaders, with the agreement that the conversation about the numbers will be exploratory rather than evaluative for the first quarter. Look for patterns: which jobs produced healthy joint margins and what they had in common, which jobs produced poor joint margins and what they had in common.

From the patterns, identify two or three operational changes that would lift the joint number. Implement them. Continue measuring. After two quarters of exploratory measurement, formalize the shared scoreboard as part of the regular leadership review of both functions, with explicit accountability and explicit linkage to the function leaders' performance evaluations.

The first quarter is uncomfortable. The second quarter is informative. By the third quarter, both functions have internalized the joint accountability and the conversation has fundamentally changed.

The full stack

The five articles in this cluster describe the full operational stack that the best restoration companies are building around the mitigation-to-reconstruction handoff. The handoff is the most expensive moment in the restoration economic chain. The prep standard is the document that makes the handoff designed rather than accidental. The photo and documentation discipline is what gives the handoff the data the rebuild team needs to perform. The feedback loop is what keeps the standard alive over years. And the shared scoreboard is what holds both functions accountable to the joint outcome and makes all the other artifacts work in concert.

None of this is technology. None of it requires capital. All of it requires operational seriousness sustained over years. The companies that build this stack are quietly creating one of the most durable competitive advantages available in the industry. The companies that do not are paying for the absence on every job, every quarter, every year, in a leak that does not show up as a single line item but that determines whether the company is on the operating-system side of the industry split — or the side that wakes up in 2028 wondering what happened.

This cluster is closed. The next clusters in [The Restoration Operator's Playbook](#) will go deep on AI in restoration operations, on financial operations discipline, on carrier and TPA strategy, and on the senior talent question. Each cluster builds on the others. Each contributes to the same underlying argument: the restoration industry is splitting into two groups, the split is happening on operational discipline, and the window in which the right side of the split can still be reached is open now.

The companies that read this body of work and act on it will know who they are. The rest will find out later.