JUNCTURE RESEARCH · WHITEPAPER

The Pre-MLR Pre-Check Playbook

Cut the MLR bottleneck without cutting the reviewer.

JUNE 2026

JUNCTURE.HEALTH

Why MLR is the bottleneck, and why AI made it worse

Picture the content system as a pipe. Generation is the wide end, review is the narrow end. For years the two were roughly matched. A brand team drafted at a pace its medical, legal, and regulatory reviewers could keep up with, and while the friction was real, the system was at least in equilibrium.

Generative AI widened the wide end and touched nothing else. Boston Consulting Group describes generative AI driving the marginal cost of marketing content toward near zero, enabling industrial-scale production. Inside pharma the working number teams quote, illustratively, is a five to tenfold jump in drafts produced. A brand that shipped two hundred assets a quarter now produces a thousand, and routes all thousand into a review function staffed for two hundred. The reviewers are the same people, reading the same way, holding the same accountability. You cannot ask a regulatory reviewer to be ten times less careful.

The data on the narrow end is sobering on its own, before any AI multiplier. Veeva reports that most biopharmas take about three weeks, on average, to deliver new content to market. Mid and large pharma teams report MLR cycles that can stretch 50 to 60 days per content piece under current workflows. A large fraction of that time is consumed not by hard judgment but by mechanical rework: a missing ISI, an efficacy claim with no fair balance nearby, a reference that does not support the sentence it is attached to.

AI also changed the texture of the work, not just the volume. Ten human-written drafts fail in ten familiar ways a reviewer learns to spot. A thousand machine-assisted variants fail in a thousand subtle ways: a fair-balance statement that drifted a clause during a rewrite, an indication phrased a hair beyond the label, a reference that no longer supports the claim it is pinned to. The defect rate per asset may even be lower. The defect surface across the portfolio is far larger. AI did not just give MLR more to read. It gave MLR more ways to be wrong to catch.

This matters more in the current enforcement climate. In 2025 the FDA closed the year with 74 letters to drug and biologic manufacturers, 42 of them aimed at direct-to-consumer television ads, and the dominant theme across the wave was not

exotic. It was omission and minimisation of risk information. Those are exactly the mechanical failures a queue under load is most likely to let slip.

The triage principle: what to automate, and what to never automate

The bottleneck is a matching problem, and the right response is a sorting decision. Every promotional asset carries two kinds of work. One kind is mechanical: countable, rule-bound, and identical every time, of the form "is the ISI present and complete" or "does this efficacy claim have risk information near it with comparable prominence." The other kind is judgment: contextual, novel, and accountable, of the form "can this broader population be supported at all," or "does this analogy cross a line a label clause cannot quite settle." A machine is excellent at the first and has no business owning the second.

The triage principle follows directly. **Automate the mechanical layer entirely. Automate none of the judgment, the authoring, the decision to publish, or the sign-off.** The pattern the industry has converged on is human-in-the-loop, where assistive automation handles the quantitative checks and qualified reviewers stay accountable for every final approval. A pre-check is the assistive layer. It surfaces issues. It does not decide.

The danger of getting this wrong is not abstract. When a vendor "speeds up review" by having a model pre-decide what is fine and surface only exceptions, it is not supporting the control. It is eroding it, in the one part of the system where 21 CFR Part 11 expects a defensible, attributable human decision. The reviewer is the name on the approval, the signature that means a person stood behind a regulated claim. Triage done well clears the reviewer's path. Triage done badly quietly removes the reviewer from the loop and leaves no one accountable behind the green light.

A useful test for any proposed automation: if it fails silently, who is harmed, and would a regulator accept "the system said it was fine" as the answer? If the honest answer is "a patient, and no," that step belongs to a human. Everything that fails loudly, visibly, and into a reviewer's queue is fair game for the machine.

What a deterministic pre-check actually checks

A pre-check earns its place by being narrow and inspectable. The word that matters is **deterministic**. Each check is a fixed rule with a defined input and a pass or fail output. It does not paraphrase, summarise, or hold an opinion. It asks a yes-or-no question and records the answer with its evidence. Claim extraction can be AI-assisted, because reading a chart or isolating a claim from prose is genuinely hard; the rule evaluation that follows is deterministic, so every verdict is reproducible and traceable. Described at the outcome level rather than as a buildable method, a pre-check resolves five things.

Claims, from text and image both. A claim is any statement that asserts something about the product: efficacy, safety, comparative, mechanism. The hard half is the image. A bar chart that shows a benefit, a callout box, a patient photo with a superimposed result, a visual that implies complete symptom resolution: the FDA treats all of these as claims, and its 2025 letters repeatedly cited misleading efficacy conveyed through visual representations, including implied complete resolution of symptoms. A check that reads only the body copy is checking half the asset.

Fair balance against efficacy claims. Fair balance is the requirement that risk information appear with prominence and readability reasonably comparable to benefit information. The standard comes straight from 21 CFR 202.1, which requires a fair balance of benefit and risk and prohibits false or misleading claims. For every efficacy claim found, the outcome is a verdict on whether corresponding risk information is present, near the claim, and carried with comparable weight. A benefit on page one and the risk language buried on page four fails, and it fails the same way every time.

ISI presence and completeness. Important Safety Information is the concise statement of the product's most significant risks, contraindications, warnings, and precautions. Every branded efficacy asset needs it. The outcome is binary at its base, present or absent, then graded: does the ISI carry the contraindications and any boxed warning, or has a section been dropped. Omission of risk information was the single most cited violation across the 2025 enforcement wave, which makes a present-and-complete ISI verdict the highest-value line a pre-check returns.

On-label indication match. An advertisement may recommend a drug only for the uses in its approved labeling. 21 CFR 202.1 ties promotion to uses for which the drug is generally recognised as safe and effective. The outcome compares each indication or population the asset implies against the approved indication statement, and flags any claim that reaches beyond the indication, names an unapproved population, or implies a use the label does not carry, with the specific clause it failed against.

Reference substantiation. Every claim that rests on data needs a reference, and the reference has to actually support the claim. The outcome links each substantiated claim to its cited source and grades the match. A claim with no reference fails. A claim whose reference says something narrower than the claim fails. The reviewer sees the gap, not a green light.

Across all five, one property does the heavy lifting: **every verdict cites the controlling clause.** A pre-check that says "this asset has problems" is nearly useless. A pre-check that says "the page-two headline has no risk information within the fold, which fails fair balance, and here is the label section the claim must align to" is something a reviewer acts on in seconds. The citation is also what makes the run auditable, which leads to the guardrails.

The hard guardrails

The five checks describe capability. The guardrails describe restraint, and in a regulated setting the restraint is the product. Four lines a pre-check must not cross.

It never writes a claim. The pre-check evaluates the claims already in the asset. It does not draft a better headline, rewrite the ISI, or propose a compliant version of a failed claim. The moment a system generates promotional language, that language is itself a claim that must be reviewed, and a tool that both writes and clears its own writing has no independent check. The pre-check reads, scores, and cites. Composing the fix is the human's job, on purpose.

It never auto-publishes. A pass is not an approval. It is a clean triage result that moves the asset to human review faster. Nothing the pre-check does releases

content to a channel. Given the volume of 2025 enforcement and that the FDA in some letters took an aggressive view of who is responsible for a promotional communication, no automated gate should be the thing that puts a message in front of a patient. A human approves. The system never does.

It backs the reviewer with a Part 11-supporting trail. When the reviewer signs off, the system's job is to make that sign-off defensible. 21 CFR Part 11 requires secure, computer-generated, time-stamped audit trails capturing who did what and when, which cannot be edited by any user, paired with attributable electronic signatures. A pre-check that supports Part 11 records every claim it extracted, every rule it ran, every verdict it returned, and the reviewer's e-signature on the decision, in an immutable trail tied to a unique identity.

The vendor is Part 11-supporting, not compliant. This distinction is deliberate and worth stating plainly to procurement. The tool provides the technical controls: a time-stamped, tamper-evident audit trail, e-signature sign-off, access control. It is the customer who validates the system under their own standard operating procedures, runs it inside their quality framework, and owns the compliance posture. A vendor does not get to declare a customer's process compliant or validated. The vendor supplies the controls; the customer validates the use. Any tool that markets itself as "Part 11 compliant" or "validated" out of the box is selling a claim it cannot make.

Reuse scoring: review only what changed

There is a second number a reviewer has rarely had, and it may matter as much as the checks. Most "new" assets are mostly recombinations of already-approved language. A reuse score breaks an asset into blocks, matches each back to the approved modules it draws from, and tags every block exact-match, light-edit, or genuinely new. If the system can tell the reviewer that an asset is eighty percent previously approved claims with two new sentences, the reviewer reads two sentences, not the whole thing, and reads them knowing exactly where the novelty is.

This is where the modular content discipline pays off, and the published numbers are real. Veeva reports that for one customer a modular approach is **helping reduce time to market by more than 50%, stemming from an increase in content reuse by 40% and a decrease in approval times by 30%**, with another company **achieving 75% of approvals in a single review cycle, cutting costs by 19%**. Companies that have optimised their MLR workflows around reuse and tiered review have seen a **57% reduction in review cycle times and a 55% drop in time spent in review meetings**.

The point for the playbook is not the specific percentage. It is the mechanism. Reuse is not a vanity metric. It is the lever that lets a small approved core feed a large output safely, because the reviewer's attention lands on the delta instead of the whole. Velocity bought by reuse is velocity you can defend. Velocity bought by skipping checks is a deviation waiting to be written up.

A worked example: Varigel

Take a fictional brand, Varigel, approved for one narrow indication with a known contraindication for patients on a common comorbidity medication. A brand team uses AI to produce sixty pieces for a congress push: emails, rep-triggered messages, social variants, localised cuts. One of them, a new digital sales aid, runs through the pre-check.

Claim extraction finds eleven claims. Nine are text. Two are visual: a bar chart showing response rate, and a callout reading "lasting relief." On-label match passes for ten of the eleven but flags one, a subhead that implies use in a broader patient population than the approved indication, and cites the exact indication clause it exceeds. Fair balance flags the response-rate chart, an efficacy claim with no risk information on the same screen, failing the prominence-and-proximity test. ISI presence passes, but the completeness grade flags that the contraindication line is absent from the ISI block. Reference substantiation links the "lasting relief" callout to its cited study and flags it: the study reports a defined duration, and "lasting" overstates it. The reuse score reads 78% previously approved, with the new subhead and the callout tagged as the only genuinely new copy.

The verdict is four findings, each with its claim, its rule, and its citation, plus a reuse score that points the reviewer straight at the two new blocks. The pre-check writes nothing. It publishes nothing. It logs the whole run, with the reviewer's eventual e-signature, in a Part 11-supporting trail. The reviewer spends the meeting on the one genuinely substantive question, whether the broader population can be supported at all, instead of the forty minutes it would have taken to find the missing contraindication line by hand.

Now scale that across the full sixty-asset push. Without a pre-check, all sixty hit MLR and the reviewer hunts for the usual scatter of mechanical breaks across the lot, hours of careful reading where any miss is a regulatory event. With a pre-check, the breaks are caught at the gate, named by defect, and bounced back to the creators before review. The clean assets arrive each carrying a reuse score, the near-total recombinations flagged as low-novelty. Same reviewer, same standard, a fraction of the load. The bottleneck did not get a faster reviewer. It got a smaller, cleaner queue.

Implementation checklist

A pre-check is a process change as much as a tool. The following sequence stands one up without disrupting the MLR system of record.

- **Pick one brand and one asset type to start.** A single brand with a stable label and one high-volume asset type (an HCP email, a rep sales aid) gives a clean baseline and a fast learning loop. Do not boil the portfolio.
- **Codify the label as the source of truth.** Load the approved label and reference library, and write down the rules: fair balance thresholds, the canonical ISI block, the approved indication and population, and the per-market and per-product tolerances. Rules are tied to the approved label and should be configurable, so the team can encode its own interpretations and risk tolerances.
- **Place the gate before review, not inside it.** The pre-check runs the moment an asset is finished and refuses to forward anything that fails a mechanical check. The reviewer should never again open a mechanically broken asset.

- **Keep the reviewer accountable and the sign-off human.** Confirm that the tool never writes a claim, never publishes, and never substitutes for the decision. The named reviewer signs off with an e-signature, backed by a Part 11-supporting audit trail.
- **Turn on reuse scoring early.** A reuse score on every asset is what concentrates review on change and makes the modular discipline pay. Track it from day one.
- **Validate under your own SOPs.** The vendor provides the controls. Your quality function validates the system for Part 11 use inside your framework and owns the compliance posture. Build that validation into the rollout, not after it.
- **Measure the right number.** Stop celebrating generation speed. Track cycle time after the gate, days from finished asset to approved asset, alongside pre-check catch-rate before MLR, off-label drift caught, and how much output is recombination of an approved core. Those are the numbers that show the bottleneck moving.
- **Then expand.** Add asset types, then markets, then brands, carrying the rule set and the reuse library forward. The cost of each new addition falls as the approved core grows.

How a pre-check fits alongside your MLR system of record

A pre-check is a **layer, not a replacement.** It complements the MLR system of record you already run, Veeva PromoMats or Vault, Aprimo, Zinc, AEM, and does not touch the review workflow, the approved-asset store, or the accountable decision. The two halves of the answer have historically lived in different tools owned by different teams: the content-generation stack makes more, the MLR workflow reviews and stores, and nothing sits between them to check an asset against the approved message before a human spends judgment on it. The pre-check fills that gap and then hands off cleanly.

This is where Juncture fits, named here as one option among several so the playbook stays honest about the category. Juncture is a pre-MLR pre-check and Share-of-

Answer layer. Before review, it pre-checks every asset against the label and reference library, flags fair-balance, ISI, on-label, and reference breaks with the verdict tied to a clause, and hands the reviewer a checked asset plus a reuse score instead of a raw pile. It runs on Azure OpenAI for inference, customer content is never used to train models, and it works on promotional content, the approved label, and public HCP-style questions, with no PHI or PII required. Reviewers consume its pre-check findings and the accountable MLR decision and final sign-off stay with your people and your system. It backs the reviewer with a Part 11-supporting audit trail and e-signature; you validate it under your SOPs.

What is available today is deliberately scoped: claims-library and approved-module ingestion, label and source ingestion, manual asset upload, SSO via Microsoft Entra, and an audit-ready report you attach to a PromoMats or MLR submission. A Veeva Vault connector, a DAM pull, and a public API are on the roadmap, not shipped, and SOC 2 is in progress. There are no published Juncture customer case studies or quantified client results, and this paper invents none. Impact is measured directly: pre-check catch-rate before MLR, MLR cycle-time, off-label drift caught, and Share-of-Answer change run over run across ChatGPT, Gemini, Perplexity, Google AI Overviews, and Claude. A pilot is small on purpose: bring one brand, one asset, and a public question set, and stand it up in weeks.

The switching cost is low because nothing about your system of record changes. The pre-check clears the asset before MLR; your MLR workflow and your approved-content store carry on exactly as they do today. That is the whole design. Relieve the constraint by changing what arrives at review, keep the accountable human decision precisely where it belongs, and let an approved, reusable core do the heavy lifting on volume.

The takeaway

The MLR bottleneck is not a staffing problem and not a speed problem. It is a sequencing problem: teams ask their most accountable people to do the cheap mechanical work first and the irreplaceable judgment work last, on ten times the volume. Move the mechanical work upstream into a deterministic pre-check, hand the reviewer a checked asset and a reuse score, keep the authoring, the decision,

and the sign-off human, and measure cycle time after the gate. The queue empties from the front, the reviewer stays the control, and the content AI let you make becomes content you can actually ship.

Sources

1. Boston Consulting Group, "GenAI's Influence on Future Creativity in Marketing," 2024. bcg.com
2. Veeva, "Getting Started With Modular Content in Pharma" (most biopharmas take about three weeks to deliver new content to market; cites the Veeva Pulse Content Metrics Report). veeva.com
3. Veeva, "Modular Content Powers Omnichannel Engagement at Speed and Scale" (time to market more than 50% lower, content reuse up 40%, approval times down 30%, 75% single-cycle approvals, 19% cost reduction). veeva.com
4. Veeva, "Future-Proof Your MLR Review and Approval Process" (57% reduction in review cycle times, 55% drop in review-meeting time). veeva.com
5. Indegene, "Is the MLR Review Process Slowing Pharma Marketing?" (MLR cycles stretch 50 to 60 days per content piece). indegene.com
6. Indegene, "Future of MLR Review in Pharma: AI and MLR Operations" (human-in-the-loop, assistive automation with accountable reviewers). indegene.com
7. King & Spalding, "2025 Year in Review: FDA Drug and Device Advertising and Promotion Enforcement" (74 letters in 2025, 42 aimed at DTC TV ads; misleading visual efficacy). kslaw.com
8. Latham & Watkins, "FDA Begins Crackdown on Direct-to-Consumer Pharmaceutical Advertising," 2025 (FDA's view on responsibility for a promotional communication). lw.com
9. Legal Information Institute, "21 CFR 202.1, Prescription-drug advertisements" (fair balance; false or misleading claims). law.cornell.edu

9. Electronic Code of Federal Regulations, "21 CFR 202.1" (promotion tied to approved uses) and "21 CFR Part 11, Electronic Records; Electronic Signatures" (time-stamped, non-editable audit trails and attributable e-signatures). ecfr.gov 202.1 · ecfr.gov Part 11