Frank Barilone Jr., MBA, CSM
Lead Product Manager – VP, Wells Fargo Home Lending Servicing

Feedback on MISMO FRAME v.01: What Works, What Is Missing, and What Needs to Happen Next

---

Where I Am Coming From

I have spent the past several years remediating OCC enforcement actions at Wells Fargo Home Lending, managing third-party oversight programs across mortgage servicing vendors, and building AI governance systems for financial services organizations through Barinhall Consulting. I am not writing this as a policy commentator. I am writing it as someone who will actually have to implement FRAME and who has already run into most of the operational problems the framework is trying to solve.

These comments are direct. That is intentional. MISMO asked for industry feedback, and the most useful feedback I can give is the kind that names the gaps clearly rather than burying them in diplomatic framing. I want FRAME to succeed. That is why I am being specific about where version 1.0 needs work.

What FRAME Gets Right

A few things first, because they matter.

The risk-tiered architecture is correct. Treating a loss mitigation recommendation engine and an internal document OCR tool as equivalent governance problems is how organizations get paralyzed. Calibrating documentation and oversight requirements to actual consumer impact and regulatory exposure is the right design choice, and FRAME does it well.

The Governance Policy's foundational principles are the best part of this entire framework. Especially these two: vendor systems are the organization's responsibility, and accountability is individual and named. No equivocation on either point. Those principles need to survive every future revision of this document intact.

I also want to call out the candor in the Getting Started Guide specifically. Acknowledging in the published framework that vendor documentation is often incomplete and that fair lending evidence for vendor AI is often missing is not the kind of admission you usually see in an industry standards document. Most frameworks present requirements as if the environment is cooperative by default. FRAME does not pretend that. That honesty is a genuine strength.

The GenAI requirements in Section 7 of the Governance Policy are ahead of where most frameworks are right now. The prohibited uses list, the prompt governance requirements, and the requirement to maintain the ability to reconstruct any system prompt in effect at any point in time for examination and litigation purposes: these are the right requirements. They reflect how these systems actually fail in regulated environments, not just how they might fail in theory.

And the regulatory mapping is genuinely useful. Having a single document that connects ECOA, FCRA, SR 11-7, OCC 2013-29, Reg X, SCRA, HMDA, and both GSE guides to specific AI system implications, by risk tier, is something practitioners have needed for a long time.

The Six Gaps

1. The Vendor Side Has No Obligations

This is the central problem with v.01, and it runs through everything else on this list.

FRAME tells servicers they are responsible for the performance, compliance, and governance of vendor AI systems, and that this responsibility cannot be contractually transferred. That is correct. Then it requires servicers to document training data, performance metrics, model limitations, adverse action explainability, regulatory testing results, result reproducibility, and proxy variable analysis for every AI system, including vendor-provided ones. For a vendor black-box model, every single one of those fields depends on the vendor choosing to provide the information.

The Getting Started Guide's answer to this is: treat vendor documentation gaps as normal and try to close them at contract renewal. I understand why that answer is in there. It is pragmatic and it reflects reality. But it creates a situation where a servicer can follow FRAME in complete good faith, document every gap, engage every vendor, and still present an examiner with a System Risk Assessment for a High tier loss mitigation engine that is essentially a list of things the vendor would not tell them.

> *FRAME currently has a servicer-side obligation and a vendor-side polite request. Those are not the same thing.*

What needs to happen here is a companion Vendor AI Disclosure Standard, developed in parallel with the AI Certification program, that establishes what information vendors of

Critical and High tier mortgage AI systems are expected to disclose. MISMO cannot impose obligations on vendors directly. But it can make vendor AI Certification meaningful enough that servicers demand it during procurement, which creates the market pressure that a voluntary disclosure standard cannot.

2. Section 8.2 Is a List, Not a Tool

The required contract provisions in Section 8.2 are the right provisions. Advance notice of model changes, audit rights, incident notification, cooperation with regulatory examinations, prohibition on using borrower data for model training without consent: these all belong in every AI vendor agreement for Critical and High tier systems. Most current mortgage tech agreements do not contain them.

The problem is that a list of required provisions without model contract language is not very useful for most of the organizations FRAME is designed to serve. An independent mortgage bank or mid-size non-bank servicer does not have a technology transactions attorney who specializes in AI governance sitting in their legal department. They have a general counsel or outside counsel who handles standard MSA reviews. Handing that person a list of required provisions and telling them to negotiate is not the same as handing them draft language.

There is also a placeholder in the current draft that needs to be resolved: the advance notice provision for version sunset and end-of-life support reads 'minimum XX days for XXX.' That placeholder should not survive into the next version. My recommendation is tiered notice periods by risk tier: 60 days for Critical systems, 30 days for High, 15 days for Medium. For incident notification, the standard should specify what counts as a notifiable event and a maximum 72-hour notification window.

MISMO should publish model contract language for the Section 8.2 provisions as a separate work product. That is one of the most practically useful things this framework can generate for smaller organizations.

3. Servicing Risk Is Under-Addressed

The framework classifies servicing systems with Reg X implications in the High tier and correctly notes that servicing AI carries significant regulatory risk that is often underappreciated. Both of those statements are true. The gap is that the current artifacts do not give servicers any additional guidance for navigating that risk environment.

Origination AI and servicing AI are different problems. Origination AI risk is primarily about credit decision fairness, adverse action accuracy, and HMDA data integrity. Servicing AI risk adds timing, notice, and process obligations that are specified by statute with a precision that origination does not face in the same way. A loss mitigation recommendation engine that misroutes a borrower does not just create fair lending exposure. It can produce a Reg X violation with hard deadline consequences. A payment processing AI that misclassifies a payment can trigger servicemember protection liability under SCRA. A borrower communication chatbot that gives incorrect

forbearance information can create consumer protection exposure that lives independently of any model bias question.

The System Risk Assessment form needs a Servicing AI Annex for use with High tier servicing systems. It should require Reg X timeline compliance testing methodology, SCRA eligibility identification accuracy testing protocol, loss mitigation waterfall adherence validation, and a borrower communication compliance review process. MISMO's Servicing Community of Practice should co-author this annex. The operational expertise required to write it correctly does not live in a general AI governance working group.

4. The Deferred Artifacts Leave a Real Gap Right Now

The System Performance and Fairness Testing File and the Incident and Change Management Log are deferred until a later stage. I understand the decision. Getting the three core artifacts out was the priority. But the deferral creates a specific problem for organizations that need to implement FRAME before those artifacts are available.

The Governance Policy establishes escalation triggers and monitoring cadences that reference documentation the deferred artifacts are supposed to contain. Section 9.2 requires immediate escalation for a material vendor model change that was not contractually disclosed in advance. Where does that escalation get documented? The Incident and Change Management Log. Which does not exist yet.

Similarly, the Governance Policy requires pre-deployment fair lending testing for Critical and High tier systems, with specific documentation requirements. The standardized format for that documentation is the System Performance and Fairness Testing File. Which is also deferred.

MISMO needs to publish interim guidance that tells organizations what minimally acceptable documentation looks like for both artifacts in the period before the standards are released. It does not need to be a full artifact. A two-page guidance note specifying required elements and acceptable formats would be sufficient. There should also be a clear published timeline for when the deferred artifacts will be available.

5. GenAI Foundation Model Changes Are a Real Problem Without a Real Solution

Section 7.4 requires organizations using third-party foundation models to maintain contractual provisions requiring advance notice of model version changes, and to conduct regression testing before activating new versions for Critical and High tier systems. That requirement is correct. Here is the operational reality.

Foundation model providers update continuously. Behavioral changes are sometimes disclosed in release notes that may or may not reach the compliance or technology risk functions at a servicer before the vendor deploys the updated version. Many mortgage technology vendors building on commercial LLMs are not pinning specific model versions in production. They are letting the underlying model update automatically

because that is the path of least resistance and their agreements with the foundation model provider do not require them to do otherwise.

A servicer whose vendor deploys a GenAI-powered servicing chatbot or document processing application has essentially no visibility into whether the underlying model changed last week, what changed, or whether any behavioral change affected compliance-relevant outputs. FRAME's requirement that the servicer maintain contractual provisions requiring advance notice of these changes is the right requirement. It does not solve the problem of foundation model providers not providing that notice in the first place.

Two specific things need to happen. First, MISMO should engage foundation model providers directly through the AI Certification program to establish a Mortgage Industry AI Disclosure Protocol for material behavioral changes in foundation models used in mortgage applications. Second, in the interim, FRAME should explicitly recommend that servicers require their GenAI vendors to pin specific model versions for production use and prohibit automatic foundation model updates without prior servicer notification and impact assessment. This needs to be stated in the framework as a requirement, not left as something servicers have to figure out on their own.

6. The System Risk Assessment Has No Protocol for When Vendors Will Not Cooperate

The System Risk Assessment is a well-designed form. It asks for the right information. The problem is what happens when the vendor will not provide it.

Right now, the form has no structured mechanism for documenting vendor non-disclosure. There is no field for 'vendor was asked and declined.' There is no guidance on what constitutes an acceptable compensating control when documentation cannot be obtained. There is no threshold that triggers an escalation or a deployment hold when the gaps are too significant to work around.

The Getting Started Guide tells servicers to capture what exists, document what is missing, and use renewals to close the gaps. From a governance maturation standpoint, that is reasonable advice. From an examination standpoint, it is not sufficient on its own. An examiner looking at a completed System Risk Assessment for a High tier loss mitigation AI with multiple blank fields and no vendor documentation may characterize that as a governance failure rather than a documented gap. FRAME should help servicers understand and navigate that distinction, and give them the tools to make the distinction legible to an examiner.

The fix here is a Vendor Documentation Status section added to the System Risk Assessment. It should include a status field for each documentation category: Provided, Requested and Pending, Requested and Declined, or Not Applicable. It should require the date of request, vendor contact, and any compensating control documented in lieu of vendor disclosure. For Critical and High tier systems where the vendor has declined to provide model documentation or fairness testing results, the form should require a documented escalation path and a remediation timeline. That creates an audit trail that demonstrates governance discipline even when vendor cooperation is incomplete.

A Few Additional Things Worth Flagging

OCC Bulletin 2026-13 is not in the current regulatory reference list. Given the overlap with SR 11-7 and the OCC's examination focus on AI governance in supervised institutions, it belongs there. I would recommend confirming alignment with 2026-13 before the next version is published and adding it to Appendix A in both the Overview and the Governance Policy template.

Fannie Mae LL-2026-04 becomes effective August 6, 2026. Organizations that begin FRAME implementation today and follow the Getting Started Guide's recommended 30-day quick start timeline may not complete even Phase 1 before that deadline arrives. The Getting Started Guide should flag this explicitly and identify which FRAME artifacts are most directly responsive to LL-2026-04's requirements.

The AI System Inventory template is missing two columns that would make it operationally useful over time rather than just at initial implementation: vendor contract status (active, renewal pending, expired) and date of last vendor governance review. Without those fields, the inventory is a point-in-time catalog rather than a living vendor management tool.

On adverse action reason code validation for vendor models: the Governance Policy correctly prohibits relying solely on vendor-generated reason codes without verifying accuracy against the vendor model's actual outputs. For large servicers running vendor models at scale, manual sampling is the only method available when vendors do not provide model internals. FRAME should acknowledge that constraint and provide guidance on what a defensible sampling-based compensating control looks like in that environment. Otherwise, servicers are left trying to satisfy a requirement with no practical guidance on how to do so without vendor cooperation.

Where This Needs to Go

The core issue running through all six of these gaps is the same: FRAME has built a rigorous, well-designed governance structure for the demand side of the AI relationship and has not yet built the corresponding structure for the supply side. The Certification program is where that gets addressed. That makes the Certification program the most important thing MISMO is working on right now, not a follow-on to FRAME but the mechanism through which FRAME becomes fully operational.

I am aware that creating vendor-side obligations through a voluntary industry framework is harder than creating lender-side governance requirements. Vendors are not subject to OCC examination. They do not have GSE counterparty agreements that create a 'must do' rather than 'should do' posture. The Certification program has to create that pressure through market mechanisms rather than regulatory ones.

That means the Certification program needs to be designed so that servicers, particularly large ones, make MIMO AI Certification a vendor selection criterion. If the five largest non-bank servicers and the largest bank servicers require FRAME-aligned AI governance certification as a condition of preferred vendor status, the market will respond. MISMO is in a position to facilitate that alignment. It should.

The feedback above is offered in the spirit the framework deserves. FRAME is the right work. These are the places where the next version can be stronger. I am happy to contribute further through the AI Community of Practice or any other channel MISMO designates for FRAME development.

Frank Barilone Jr., MBA, CSM
Lead Product Manager - VP, Wells Fargo Home Lending Servicing