Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 7 October 2026

T. Adebayo
O. Apalowo
F. Makanjuola
Veridom Ltd

5 April 2026

OMP Domain Profile: AI Governance and Accountability Evidence for US Housing Finance Under FHFA Bulletin 2025-16 and GSE AI/ML Model Risk Governance

draft-veridom-omp-fhfa-00

Abstract

This document defines a domain profile of the Operating Model Protocol (OMP) for AI and machine learning (ML) systems deployed in US housing finance contexts subject to the Federal Housing Finance Agency (FHFA) Bulletin 2025-16 (effective March 3, 2026), which establishes a comprehensive AI governance framework for Fannie Mae, Freddie Mac, and the Federal Home Loan Banks (the GSEs), requiring transparency, accountability, and ethical stewardship for AI/ML systems used in housing finance decisions.

The profile -- designated HomeMark -- specifies how OMP's deterministic routing invariant, Watchtower enforcement framework, and three-layer cryptographic integrity architecture satisfy the AI governance evidence requirements of FHFA Bulletin 2025-16, including per-decision accountability, named individual responsibility, model risk governance documentation, fair lending evidence, and representation and warranty compliance for mortgage origination, credit decisioning, property valuation, and loan servicing.

The OMP core specification is defined in the Operating Model Protocol Internet-Draft (draft-veridom-omp).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Copyright Notice

Table of Contents

1. Introduction

AI and machine learning systems are now foundational to US housing finance operations: in automated underwriting systems (AUS) for mortgage origination, automated valuation models (AVMs) for property assessment, loss mitigation decisioning in loan servicing, and loan acquisition models in the secondary market. The GSEs operate at national scale -- Fannie Mae and Freddie Mac collectively support the majority of US mortgage originations -- meaning that AI/ML governance failures have systemic implications for housing access, fair lending, and financial stability.

FHFA Bulletin 2025-16 [FHFA-2025-16] (effective March 3, 2026) establishes four governance pillars: transparency (GSEs must explain AI/ML decisions to regulators, counterparties, and borrowers at the individual loan level); accountability (named individuals must bear documented responsibility for AI/ML outcomes at scale); ethical stewardship (AI/ML systems must not produce discriminatory outcomes inconsistent with the GSEs' statutory mission); and model risk governance (AI/ML systems must be subject to rigorous MRM frameworks including decision-level reconstructability).

These requirements converge on a per-decision accountability problem that OMP [I-D.veridom-omp] is specifically designed to address: for any individual mortgage credit decision, property valuation, or servicing action influenced by AI/ML, the entity must demonstrate what the AI/ML recommended, what data it used, which named individual bore accountability, and whether the record has remained intact.

This document defines the HomeMark profile: the domain-specific instantiation of OMP for FHFA-regulated housing finance AI/ML deployments. HomeMark denotes that every AI/ML-assisted housing finance decision is cryptographically marked against the entity's FHFA Bulletin 2025-16 obligations, producing a tamper-evident accountability record at the loan level.

Related OMP domain profiles include the Employment ADS profile [I-D.veridom-omp-employ] and the EU AI Act Article 12 profile [I-D.veridom-omp-euaia]. Audit Trace payloads are canonicalized per [RFC8785]. The OMP specification is also archived at [ZENODO-OMP].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] [RFC8174].

2. Terminology

This document uses the terminology defined in [I-D.veridom-omp]. In addition:

* Government-Sponsored Enterprise (GSE): Fannie Mae, Freddie Mac, or a Federal Home Loan Bank, as regulated by FHFA under the Housing and Economic Recovery Act of 2008.

* Automated Underwriting System (AUS): A GSE-operated or GSE-approved AI/ML system that evaluates mortgage applications and provides a credit recommendation (Approve/Eligible, Refer, Refer with Caution, or Ineligible). Includes Fannie Mae Desktop Underwriter (DU) and Freddie Mac Loan Product Advisor (LPA).

* Automated Valuation Model (AVM): An AI/ML system that generates an estimate of a property's market value based on comparable sales data, property characteristics, and market conditions.

* Consequential Housing Finance Decision: An AI/ML-assisted decision that directly affects a borrower's mortgage application status, loan terms, property valuation, or loan servicing outcome. Subject to the HomeMark Invariant.

* Responsible Individual (RI): The named individual within a GSE, lender, servicer, or counterparty who bears documented accountability for an AI/ML-assisted housing finance decision. In OMP terms, the Named Accountable Officer for ASSISTED and ESCALATED interactions.

* Representation and Warranty (R&W): The representations and warranties made by mortgage originators and sellers to the GSEs regarding loan quality, eligibility, and compliance. Where AI/ML contributed to a loan-level decision, R&W obligations require the ability to demonstrate that the AI/ML operated correctly and consistently with applicable guidelines.

* Fair Lending Flag: A field indicating that the AI/ML recommendation involves a borrower demographic profile or geographic area identified in the entity's fair lending analysis as requiring heightened review for potential disparate impact under ECOA [ECOA] or the Fair Housing Act [FHA-1968].

* HomeMark Invariant: The two-property invariant defined in Section 7: every Consequential Housing Finance Decision generates a sealed HomeMark Audit Trace independently verifiable by FHFA examiners, counterparties, and auditors.

3. FHFA AI Governance Framework Analysis

3.1. FHFA Bulletin 2025-16

FHFA Bulletin 2025-16 requires transparency (individual loan-level documentation explaining AI/ML decisions, contemporaneous not retrospective), accountability (named Responsible Individuals with documented responsibility for AI/ML system governance and decision outcomes), ethical stewardship (per-decision fair lending monitoring and disparate impact assessment), and model risk governance (decision- level reconstructability, ongoing performance monitoring, and human oversight at defined thresholds).

3.2. GSE AI/ML Model Risk Governance

GSE MRG frameworks, informed by Bulletin 2025-16 and SR 11-7 [SR-11-7], require decision- level reconstructability (for any loan-level decision, the entity must reconstruct the model's input data, version, and output consistent with the specific loan record); ongoing monitoring for performance degradation, distributional shift, and fair lending risk; and human oversight documentation at defined thresholds.

3.3. Fair Lending Obligations and Disparate Impact

The GSEs operate under ECOA and the Fair Housing Act, prohibiting both intentional discrimination and AI/ML practices producing unjustified disparate impact. FHFA Bulletin 2025-16 requires GSEs to assess and document disparate impact in AI/ML- assisted housing finance decisions. The per-decision HomeMark Audit Trace provides the loan-level evidence fair lending examinations require: what the AI/ML recommended, what data it used, whether a fair lending flag was triggered, and what human oversight was applied.

3.4. Representation and Warranty Framework

Where an AI/ML system contributed to loan origination or eligibility determination, the seller's R&W obligations require the ability to demonstrate that the AI/ML operated correctly and consistently with applicable guidelines at origination. HomeMark Audit Traces generated at origination provide this loan-level evidence: the RFC 3161 [RFC3161] timestamp proves the AI/ML recommendation was generated at origination (not reconstructed retrospectively), the interaction_hash proves input data integrity, and the ai_ml_system_version documents which AUS version was in effect.

3.5. FHFA Examination Authority

FHFA has broad examination authority over GSEs and their counterparties. FHFA examiners may request AI/ML decision process documentation, model risk governance evidence, and fair lending monitoring data at the individual loan level. The HomeMark FHFA Examination Package is designed to satisfy examiner requests within the 30-second production capability specified in this profile.

3.6. Convergent Requirements

FHFA Bulletin 2025-16, GSE MRG frameworks, ECOA/FHA obligations, and the R&W framework converge on a structure mapping to OMP's three routing states: AI/ML decisions where the RI reviewed the recommendation and bears documented accountability correspond to ASSISTED; decisions where a Fair Lending Flag triggered, confidence fell below the housing finance floor, or a model governance concern was detected correspond to ESCALATED; fully autonomous AUS-eligible transactions are permitted under AUTONOMOUS subject to Section 4.1 constraints, but HomeMark Audit Traces MUST be generated even for AUTONOMOUS routing.

4. OMP HomeMark Profile

4.1. Routing States Under This Profile

* AUTONOMOUS: Permitted for standard AUS-eligible mortgage transactions where: the AUS recommendation is Approve/Eligible; the Confidence Score meets the AUTONOMOUS threshold; no Watchtower has triggered; the Fair Lending Flag has not been set; and the loan falls within the AUS's validated operating envelope. Even under AUTONOMOUS routing, the HomeMark Audit Trace MUST be generated and sealed for every loan-level interaction, consistent with FHFA Bulletin 2025-16's transparency and reconstructability requirements.

* ASSISTED: Required where: AUS recommendation is Refer or Refer with Caution; transaction exceeds the significance threshold requiring RI review; a Fair Lending Flag is set; or a model governance concern is detected. The RI's identity, review timestamp, and decision basis are sealed in the HomeMark Audit Trace.

* ESCALATED: Triggered by: HARD_BLOCK from WT-FHFA-02, confidence failure below the housing finance decision floor (WT-FHFA-01), model performance anomaly (WT-FHFA-05), or regulatory override requirement. AI/ML recommendation MUST NOT be acted upon until the RI has reviewed and documented a disposition.

4.2. Named Accountable Officer: The Responsible Individual

The Named Accountable Officer under this profile is the Responsible Individual: the named person who bears documented accountability for the AI/ML-assisted housing finance decision. Required fields:

* ri_employee_id: stable identifier consistent throughout the relevant loan warranty period;
* ri_role: role in the AI/ML governance or decision chain (e.g., "underwriter", "credit_officer", "AUS_governance_lead");
* ri_review_timestamp: ISO 8601 UTC of the RI’s review action;
* ri_decision: one of PROCEED_WITH_AI_RECOMMENDATION, PROCEED_MODIFIED, OVERRIDE, DENY_APPLICATION, APPROVE_APPLICATION, REFER_TO_MANUAL_UNDERWRITING;
* ri_decision_basis: REQUIRED for all values other than PROCEED_WITH_AI_RECOMMENDATION.

4.3. Watchtower Definitions

4.3.1. WT-FHFA-01: Housing Finance Decision Floor Gate

Trigger: Composite Confidence Score falls below the housing finance decision floor. For AUS: a Refer or Refer with Caution recommendation signals the loan is outside AUS approval parameters.

Action: FORCE_ASSISTED. RI reviews the AI/ML recommendation before any credit action. Loan file MUST reflect the RI’s documented review.

Rationale: FHFA Bulletin 2025-16 requires human oversight of AI/ML decisions below defined confidence thresholds. An AUS Refer recommendation is itself a signal that human underwriting review is required.

4.3.2. WT-FHFA-02: Fair Lending Override Gate

Trigger: AI/ML recommendation involves a borrower demographic profile, geographic area, or loan characteristic identified in the entity’s fair lending analysis as requiring heightened review for potential disparate impact under ECOA or the Fair Housing Act.

Action: FORCE_ASSISTED for standard heightened review. HARD_BLOCK where the AI/ML recommendation conflicts with a pre-identified fair lending risk pattern in the entity’s corrective action plan.

Rationale: ECOA and the Fair Housing Act prohibit disparate impact in mortgage credit decisioning. WT-FHFA-02 ensures loan applications in identified heightened-review categories receive documented human oversight, sealed in the Audit Trace for FHFA examination.

4.3.3. WT-FHFA-03: Fair Lending Flag Gate

Trigger: Ongoing HomeMark Audit Trace monitoring identifies that the AI/ML recommendation falls within a demographic or geographic segment exhibiting an approval rate or pricing disparity above the entity’s configured fair lending alert threshold.

Action: FORCE_ASSISTED. fair_lending_flag set to true. RI review and decision basis REQUIRED.

Rationale: Continuous per-decision fair lending monitoring enables entities to identify emerging disparate impact before it reaches the threshold of a CFPB or FHFA examination finding. WT-FHFA-03 converts a periodic audit obligation into a continuous per-decision flag.

4.3.4. WT-FHFA-04: AVM / AUS Training Limitation Gate

Trigger: Property or loan characteristics match a known validation limitation of the AVM or AUS model (e.g., property type with limited comparable sales data; geographic market where the AUS was not validated; loan product feature outside the validated operating envelope).

Action: FORCE_ASSISTED. HomeMark Audit Trace records the specific training limitation triggered and the RI’s disposition.

Rationale: GSE model risk governance frameworks require entities to document model limitations and ensure decisions outside the validated envelope receive human review. WT-FHFA-04 gives this requirement structural enforcement at the per-decision level.

4.3.5. WT-FHFA-05: Model Performance Anomaly Gate

Trigger: AI/ML recommendation deviates from expected operating parameters suggesting model degradation, distributional shift, or data quality failure.

Action: FORCE_ESCALATED plus model performance anomaly alert to the entity’s model risk governance team.

Rationale: AI/ML models in housing finance can experience distributional shift as housing market conditions evolve. Early detection prevents systematic portfolio-level impact from a degraded model operating at scale.

4.3.6. WT-FHFA-06: R&W Eligibility Verification Gate

Trigger: For loans destined for GSE sale: the AI/ML recommendation or loan data presents a characteristic requiring specific verification for GSE representation and warranty compliance (e.g., loan type requiring additional documentation; data field outside AUS verified input range; characteristic identified in recent GSE quality control findings as a common R&W breach source).

Action: FORCE_ASSISTED. RI verifies eligibility before the loan proceeds to GSE sale. HomeMark Audit Trace records eligibility verification and RI confirmation.

Rationale: GSE R&W obligations require originators to represent that AUS input data was accurate. WT-FHFA-06 creates a sealed per-loan eligibility verification record supporting R&W compliance and providing evidence in any subsequent repurchase demand.

4.4. Audit Trace Schema Extensions

The following fields are REQUIRED under the HomeMark profile, in addition to core fields in [I-D.veridom-omp] Section 7:

* ri_employee_id: string, REQUIRED for Consequential Housing Finance Decisions.

* ri_role: string, REQUIRED.

* ri_review_timestamp: string, ISO 8601 UTC, REQUIRED for ASSISTED and ESCALATED.

* ri_decision: string, REQUIRED for ASSISTED and ESCALATED. One of: PROCEED_WITH_AI_RECOMMENDATION, PROCEED_MODIFIED, OVERRIDE, DENY_APPLICATION, APPROVE_APPLICATION, REFER_TO_MANUAL_UNDERWRITING.

* ri_decision_basis: string, OPTIONAL for PROCEED_WITH_AI_RECOMMENDATION; REQUIRED for all other values.

* loan_identifier: string, REQUIRED. Unique loan number or mortgage ID, enabling per-loan audit trail retrieval for FHFA examination, R&W review, and fair lending investigation.

* ai_ml_system_id: string, REQUIRED. Identifier of the AI/ML system (e.g., "DU_v11.0", "LPA_v5.2").

* ai_ml_system_version: string, REQUIRED. Version in effect at time of interaction. Critical for model risk governance and R&W compliance.

* ai_ml_recommendation: string, REQUIRED. The AI/ML output: for AUS, one of "approve_eligible", "refer", "refer_with_caution", "ineligible"; for AVM, the estimated value and confidence interval; for servicing, the recommended loss mitigation option.

* fair_lending_flag: boolean, REQUIRED. True if WT-FHFA-02 or WT-FHFA-03 triggered.

* fair_lending_basis: string, REQUIRED if fair_lending_flag is true. The demographic or geographic basis for the flag.

* housing_decision_category: string, REQUIRED. One of: "mortgage_origination", "automated_valuation", "loan_servicing", "loss_mitigation", "secondary_market_acquisition".

* rw_eligibility_verified: boolean, REQUIRED for loans destined for GSE sale. True if WT-FHFA-06 evaluated and confirmed R&W eligibility.

* fhfa_bulletin_version: string, REQUIRED. Set to "FHFA-2025-16" for deployments under Bulletin 2025-16.

* profile_version: string, REQUIRED. MUST be "VERIDOM-HOMEMARK-v1.0".

5. Representation and Warranty Evidence Architecture

The GSE R&W framework creates a retrospective evidence requirement: years after origination, lenders may face repurchase demands requiring demonstration that AI/ML was used correctly. HomeMark Audit Traces provide three specific R&W properties: contemporaneity (RFC 3161 timestamp proves the Audit Trace was generated at origination, not retrospectively); input data integrity (interaction_hash proves the AUS input data has not been altered); and AI/ML system version documentation (ai_ml_system_id and ai_ml_system_version prove which AUS version was in effect at origination).

Lenders delivering loans to GSEs SHOULD generate HomeMark Audit Traces for all AUS-assisted originations and retain them for the full R&W warranty period (typically seven years from the note date or loan payoff, whichever is later).

6. Fair Lending Evidence Package

The HomeMark profile generates per-loan fair lending evidence (each Audit Trace contains fair_lending_flag and fair_lending_basis) and aggregate fair lending evidence (the Audit Trace stream can be aggregated to compute approval rates by demographic segment, disparate impact ratios, and pricing disparities for FHFA Bulletin 2025-16 monitoring and HMDA analysis).

The Fair Lending Evidence Package for a defined loan portfolio MUST contain: all sealed HomeMark Audit Traces organised by housing_decision_category and ai_ml_system_id; aggregate approval rate data by fair lending segment; disparate impact ratio calculations; count and disposition of WT-FHFA-02 and WT-FHFA-03 triggers; chain integrity proof (SHA-256 Merkle root); and RFC 3161 TimeStampToken verification from the OMP Reference Validator [OMP-OPEN-CORE]. FHFA examiners can verify completeness and integrity without relying on the entity's reconstructed data.

7. The HomeMark Invariant

Implementations of this profile MUST satisfy the following two-property invariant:

* Property 1 (Housing finance decision accountability completeness): Every Consequential Housing Finance Decision MUST generate a sealed HomeMark Audit Trace containing: the AI/ML recommendation; the RI's identity and review timestamp where ASSISTED or ESCALATED; the RI's decision and basis where required; the Fair Lending Flag evaluation; the AI/ML system identity and version; and R&W eligibility verification where applicable.

* Property 2 (Immutable trail): The HomeMark Audit Trace MUST be sealed with the three-layer integrity architecture defined in [I-D.veridom-omp] Section 7. Any modification to any historical Audit Trace record MUST be detectable by FHFA examiners, GSE counterparties, or any third-party auditor without access to the entity's or OMP implementer's infrastructure.

An entity satisfying the HomeMark Invariant can demonstrate, for any Consequential Housing Finance Decision: the AI/ML recommendation and input data; the AI/ML system identity and version; the RI's identity and review timestamp; the RI's decision and independent basis; the

Fair Lending Flag status; R&W eligibility verification where applicable; and that the record has not been altered since sealing. This satisfies the transparency, accountability, and model risk governance requirements of FHFA Bulletin 2025-16, the fair lending examination evidence standards of ECOA and the Fair Housing Act, and the R&W compliance evidence standards of the GSE selling and servicing frameworks.

8. Security Considerations

The security considerations of [I-D.veridom-omp] apply in full.

Borrower data sensitivity: HomeMark Audit Traces contain borrower PII and financial data subject to GLBA privacy requirements. Operators MUST implement GLBA-compliant safeguards. Fair lending demographic data used in WT-FHFA-02 and WT-FHFA-03 MUST be segregated from credit decision data consistent with ECOA's prohibition on using protected characteristics in credit decisions.

AI/ML system version integrity: The ai_ml_system_version field is a critical R&W and model risk governance element. Operators MUST implement controls ensuring the version recorded matches the AUS or AVM version actually in effect at decision time. Version misrepresentation is a material R&W compliance issue.

Loan identifier uniqueness: The loan_identifier field MUST be globally unique within the operator's deployment. Duplicate identifiers would create ambiguity in per-loan evidence retrieval and undermine R&W compliance documentation.

RI identity integrity: ri_employee_id MUST reflect the individual who actually reviewed or was accountable for the AI/ML-assisted decision. Operators MUST implement technical controls preventing RI identity assignment without the relevant individual's authenticated action.

9. IANA Considerations

This document has no IANA actions.

10. References

10.1. Normative References

[I-D.veridom-omp]
Adebayo, T., Apalowo, O., and F. Makanjuola, "Operating Model Protocol (OMP): A Deterministic Decision-Enforcement Protocol with Externalized Proof-of-Integrity", Work in Progress, Internet-Draft, draft-veridom-omp-00, March 2026, <https://datatracker.ietf.org/doc/html/draft-veridom-omp-00>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3161] Adams, C., Cain, P., Pinkas, D., and R. Zuccherato, "Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP)", RFC 3161, August 2001, <https://www.rfc-editor.org/info/rfc3161>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8785] Rundgren, A., Jordan, B., and S. Erdtman, "JSON Canonicalization Scheme (JCS)", RFC 8785, June 2020, <https://www.rfc-editor.org/info/rfc8785>.

10.2. Informative References

[ECOA] U.S. Congress, "Equal Credit Opportunity Act, 15 U.S.C. 1691 et seq.", 1974.

[FHA-1968] U.S. Congress, "Fair Housing Act, 42 U.S.C. 3601 et seq.", 1968.

[FHFA-2025-16]
Federal Housing Finance Agency, "Bulletin 2025-16: Artificial Intelligence Governance Framework for the Enterprises and Federal Home Loan Banks", March 2026.

[I-D.veridom-omp-employ]
Adebayo, T., Apalowo, O., and F. Makanjuola, "OMP Domain Profile: Automated Decision Systems Accountability in Employment Under California FEHC CRC Regulations, New York City Local Law 144, and Related ADS Accountability Obligations", Work in Progress, Internet-Draft, draft-veridom-omp-employ-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-veridom-omp-employ-00>.

[I-D.veridom-omp-euaia]
Adebayo, T., Apalowo, O., and F. Makanjuola, "OMP Domain Profile: EU AI Act Article 12 Logging and Traceability Requirements for High-Risk AI System Operators", Work in Progress, Internet-Draft, draft-veridom-omp-euaia-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-veridom-omp-euaia-00>.

[OMP-OPEN-CORE]
Veridom Ltd, "OMP Open Core: Reference Validator and Schema Library", Apache 2.0, https://github.com/veridomltd/omp-open-core, 2026.

[SR-11-7] Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency, "Guidance on Model Risk Management (SR 11-7 / OCC 2011-12)", April 2011.

[ZENODO-OMP]
Adebayo, T., Apalowo, O., and F. Makanjuola, "OMP -- Operating Model Protocol: A Deterministic Routing Invariant for Tamper-Evident AI Decision Accountability in Regulated Industries", Zenodo DOI 10.5281/zenodo.19140948, March 2026.

Authors' Addresses

Tolulope Adebayo
Veridom Ltd
London
United Kingdom
Email: tolulope@veridom.io

Oluropo Apalowo
Veridom Ltd
Awka
Nigeria
Email: ropo@veridom.io

Festus Makanjuola
Veridom Ltd
Toronto
Canada
Email: festus@veridom.io