CMS-10410 (Medicaid Program; Eligibility Changes under the Affordable Care Act of 2010)

Submitted on behalf of Promise as an organizational comment to CMS-10410

Promise.

Comment

Introduction (who/why):
We submit this comment as a public-benefit technology practitioner focused on eligibility, verification, and accessible channels (web/SMS/voice). Our aim is to support CMS' objectives to streamline eligibility and renewals while reducing administrative burden and improving equity for Limited English Proficient (LEP), low-connectivity, and disabled populations.

Background (scope I'm addressing):
This comment addresses the *information collection* needed to (1) support real-time, event-based verification; (2) reduce resubmission burden through data reuse; and (3) ensure multimodal, accessible collection channels (including IVR/virtual agents) that capture the same minimum-necessary data with auditable provenance.

Analysis (burden, quality, trade-offs):
States can significantly lower error and rework by standardizing machine-readable schemas, enabling API-based verification, and letting applicants correct ambiguous data in-line. Voice channels—when designed with multilingual prompts, consent capture, and secure handoffs—extend access to households with limited broadband or literacy. The trade-off is ensuring automation does not increase erroneous denials; this is best mitigated by provenance logging, human-in-the-loop (HITL) exception handling, and model-risk monitoring.

Recommendations (specific, implementable):

1. **Accessible, Equivalent Channels (Voice/IVR & Virtual Agents).**
   Recognize voice/IVR as a first-class eligibility/renewal intake channel. Require parity with web forms (same fields, validations, and consent) and explicit logging of: channel, language, timestamp, captured consent, and staff/AI escalation outcomes.

2. **Event-Based, Real-Time Verification.**
   Encourage APIs to ingest authoritative data (e.g., wage, residency) with clear timing and frequency, plus "verify-once, reuse-everywhere" rules to prevent redundant requests across Medicaid/CHIP programs.

3. **Applicant-In-The-Loop (AITL) Corrections.**
   Permit applicants to confirm or correct low-confidence fields (e.g., income

amounts parsed from documents) via web, SMS, or IVR. Require that each corrected field stores source, timestamp, confidence score, and user attestation for audit.

4. **Provenance & Auditability.**
   In the ICR instructions/supporting statement, specify field-level provenance (source system/document ID, extraction/match method, confidence) and require exportable audit logs to support appeals and QC.

5. **Model-Risk Controls for Automated Matching.**
   Require basic quality management: false-positive/negative tracking on data matches, periodic calibration/drift checks, and documented thresholds that trigger HITL review—reported at least annually.

6. **Privacy/Security by Design.**
   Emphasize PHI minimization, role-based access, encryption in transit/at rest, retention limits, and vendor control expectations (e.g., SOC 2/StateRAMP-equivalent), without mandating a single cloud.

7. **Equity & Accessibility Measures.**
   Require multilingual prompts for prevalent languages, low-bandwidth fallbacks (SMS/voice), disability-accessible flows, and **metrics** by channel (completion rate, time-to-decision, abandonment, successful renewal rate) disaggregated where feasible.

8. **Standard Schemas & Validation.**
   Publish machine-readable data dictionaries (JSON/XSD) for eligibility/renewal fields and provide validation rules to reduce keystroke error and rework. Invite states to map their screens to these schemas to cut duplication.

Conclusion:
These changes will reduce burden, improve accuracy, and expand equitable access without sacrificing program integrity. We support continuation of CMS-10410 and recommend incorporating the channel parity, AITL, provenance, and model-risk elements above into the ICR's instructions and supporting materials so states can operationalize them consistently. Thank you for the opportunity to comment.

Attachment 1 — Accessible Channel Parity Checklist

Use this as a one-pager you can paste/upload. Check what you already meet; circle planned items + target dates.

Program / State: _____ Submitter: _____ Date: _____

A) Scope & Channel Parity

- ☐ Intake/renewal data elements are identical across Web / SMS / Voice (IVR or virtual agent) / Assisted.
- ☐ Same validations (required fields, formats, eligibility rules) across channels.
- ☐ Same consent text and attestation captured across channels (language-appropriate).
- ☐ Channel, language, and session identifiers are logged with each submission.

B) Consent & Provenance Logging

- ☐ Per field, we log: source (applicant, doc, authoritative dataset, staff), extraction method, confidence score, timestamp.
- ☐ Consent object includes method (web/sms/voice/paper), version of disclosure, and user attestation.
- ☐ Voice channel stores audio hash / transcript segment IDs tied to fields.

C) Identity, Privacy & Security

- ☐ Minimum necessary PHI collected; PII class tagged per field.
- ☐ Encryption in transit/at rest; role-based access; retention schedule applied automatically.
- ☐ Vendor controls: SOC 2 (or StateRAMP/FedRAMP-equivalent) documented.
- ☐ IP/device/session hashes logged for audit; no raw device fingerprints retained beyond policy.

D) LEP & Accessibility

- ☐ Multilingual flows for top LEP languages; plain-language prompts.
- ☐ Voice flows meet TTY/relay accommodations; pause/repeat; barge-in; escalation to human.
- ☐ Low-bandwidth SMS fallbacks for web/voice failures.

E) Model-Risk & Quality (for automated matching/extraction)

- ☐ False-positive / false-negative rates tracked by data source and channel.
- ☐ Drift monitoring with thresholds that trigger HITL review.
- ☐ Versioned matching/extraction algorithms with change logs.

F) Performance & Equity Metrics (disaggregated where feasible)

- ☐ Completion rate by channel & language
- ☐ Time-to-decision (median, p90)
- ☐ Abandonment rate by step
- ☐ IVR/Virtual agent containment target: ____% (current: ____%)
- ☐ Human handoff target: ____% (current: ____%)
- ☐ Renewal success rate by channel & LEP status

G) Governance & Evidence

- ☐ Annual accessibility/LEP review and stakeholder testing plan.
- ☐ Appeals/QC workflows can export full provenance (field-level) on request.
- ☐ Data-sharing agreements reflect **verify-once, reuse-everywhere** with consent controls.

Attachment 2 — Model-Risk & Quality Metrics

Metric	Definition	Unit	Target	Calculation	Notes
IVR Containment Rate	% of calls resolved without human	%	≥ 40%	resolved_without_handoff / total_calls	Track by language & intent
Human Handoff Rate	% routed to human	%	≤ 35%	handoffs / total_calls	Includes warm transfers
Time-to-Decision (p50/p90)	From submission to adjudication	minutes	p50 ≤ 60	median/ percentile	Split initial vs renewal
FP Match Rate	False positives in data matching	%	≤ 2%	fp / total_matches	By source (wage, PARIS, etc.)
FN Match Rate	False negatives	%	≤ 2%	fn / total_matches	Requires adjudicated ground truth
Extraction Confidence	Mean confidence for AI-extracted fields	0–1	≥ 0.90	avg(confidence)	Flag < threshold to AITL
AITL Resolution Rate	% low-confidence fields corrected/confirmed by applicant	%	≥ 80%	resolved / flagged	By channel
Appeal Overturn Rate	% decisions reversed on appeal	%	≤ 3%	overturned / total_decisions	Monitor drift/bias
Equity Gap	Completion rate gap (English vs LEP)	p.p.	≤ 5 p.p.	completion_en – completion_lep	By channel

Attachment 3 — Human-in-the-Loop Governance Example Index

Figure Ref	Simple Example	Governance Principle
Figure 1	*AI highlights "pay now" on a utility bill and suggests Crisis award.* Caseworker must still confirm the document's full context and check if any conflicting policy applies.	NLP recommendations must be confirmed by a human before award.
Figure 2	*AI pre-fills gross income from two paystubs.* Caseworker must confirm dates span a full 30-day window before proceeding.	Field extraction must be human-verified before used in eligibility logic.
Figure 3	*A SNAP award letter is detected by AI, but the expiration date is unclear.* Caseworker reviews the letter and manually enters eligibility end date.	Ambiguous extraction must defer to human confirmation.
Figure 4	*An ID image is flagged by AI as expired due to OCR.* Caseworker sees the date is actually MM/DD/YYYY format (not DD/MM/YYYY) and overrides the AI suggestion.	Overrides must be logged and permitted when confidence is low.
Figure 5	*AI generates a case note: "Application missing ID."* Staff edits the message to clarify which household member is missing ID before submission.	AI-generated messages must be editable and attributed.

Figure 6	*AI recommends "Home Energy award only" due to missing disconnection language.* Caseworker spots "service interruption" in the document and upgrades to Crisis.	AI cannot be the final authority on policy classification.
Figure 7	*AI chatbot answers: "You likely qualify if your income is below $2,000."* System includes a note: "This is an estimate. Please submit your application for a full review."	AI instructional content must include disclaimers and hand-off cues.
Figure 8	*AI suggests an outreach message to clients who haven't submitted documents within 5 days.* Staff edits the message tone and adds legal disclaimer.	Communication AI must be editable and logged.
Figure 9	*Planning AI shows a ZIP code with a 300% increase in denials.* Internal review identifies a local document ID bug—not a true policy issue.	Forecasting AI outputs must be reviewed before policy is adjusted.
Figure 10	*Fine-tuned model is deployed for fraud flagging.* Performance audit shows it overflags multi-generational households. Model is rolled back and retrained.	Fine-tuned models must retain rollback, audit trail, and fairness logs.

Attachment 4 — Bibliography Appendix

Sources

- Florida CS/SB 1680 (2024) [CS/SB 1680, L83-L91]: Establishing AI ethics and provenance standards.
- Florida DCF RFI [Florida DCF RFI, L11-L15]: AI-driven fraud detection initiatives.
- OMB Memo M-24-10 [OMB Memo M-24-10, L735-L743]: Federal guidance for AI governance.
    - OMB M-24-10, Section B.3 (AI-generated public outputs must be attributed, editable, and logged)
    - OMB M-24-10, Section B.2–B.4 (Oversight, Auditability, Human Review)
    - OMB M-24-10, Section B.5 (Forecasting systems must undergo review if informing operational decisions)
    - OMB M-24-10, Section C.3: "Updates to AI systems must preserve original governance guarantees."
- NIST AI Risk Management Framework [NIST AI Risk Management Framework, L7-L9]: Traceability and risk mitigation practices.
    - NIST AI RMF, Section 4.3 (Risk Tiers for Planning and Surveillance Use Cases)
    - NIST AI RMF, Section 4.4.4: “Versioning and documentation of retraining events”
- HHS AI Strategy [HHS AI Strategy, L5-L8]: Program integrity and trustworthy AI.
- HHS "Public Benefits and AI" [HHS "Public Benefits and AI," L705-L713, L727-L735]: AI oversight and fairness.
    - HHS “Public Benefits & AI,” p. 11–13 (Prohibition on AI-only determinations)
    - HHS “Public Benefits & AI,” p. 13–15 (Transparency and explainability in communications that affect access)
    - HHS “Public Benefits & AI,” p. 17 (Clear communication obligations for AI in benefits access)
- White House AI Bill of Rights, Principle #1 (Notice and Explanation) and Principle #3 (Human Alternatives)

- California SB 1120 [California SB 1120, L253-L261]: Mandated human oversight in health and benefits eligibility.
    - California SB 1120, Sec. 2 (Human review required for algorithmically influenced eligibility decisions
    - California SB 1120, Sec. 2(d) (Instructional systems may only assist—not predict—outcomes)
    - California SB 1120, Sec. 2(f) (Public service messages generated by AI must disclose origin and allow intervention)
- Colorado SB 24-205 [Colorado SB 24-205, L232-L240]: High-risk AI obligations, bias prevention.
    - Colorado SB 24-205: "Material change in model weights must be reviewed at the same level as new model deployment."
    - Colorado SB 24-205, Sec. 4(c) (Audit trail and impact classification of high-risk AI)
    - Colorado SB 24-205, Sec. 5 (Generative content must be auditable and subject to correction mechanisms)
    - Colorado SB 24-205, Sec. 6 (Transparency and traceability for instructional or communicative AI)
    - Colorado SB 24-205, Section 7 (Risk tiering of non-determinative systems)
- WBUR On Point AI welfare fraud analysis [WBUR On Point, L115-L123]: Risks of AI-driven benefit denials.
- Los Angeles County AI chatbot pilot [Route Fifty, L242-L250]: AI documentation best practices.
- FEMA Eligibility Decisions [FEMA Eligibility Decisions, L338-L342]: Identity verification standards for disaster aid.
    - FEMA Policy 104-009-18 (Fraud detection and disaster analytics must include human adjudication in all final actions)
- R Street Institute analysis [R Street Institute, L82-L89]: Balanced AI regulatory trends in state policies.