




# Use Case Best Practice Guide

Tiger Team on AI &  
Medicaid

Use Case 2: Adjudication

*V1.0*





# TABLE OF CONTENTS

3	Intro: Community Engagement Requirements
4	What Does This Guide Include?
5	Who Is This Guide For?
5	Use Case 2 Example Workflow
6	Best Practices by Principle Area
14	AI Ideas
17	Appendix 1: Examples
20	Appendix 2: Methods
21	Appendix 3: Thank you & Contributors

*Page numbers refer to this use case guide only*

# AI for Medicaid Adjudication

*For Developers & Implementers*

## Intro: Community Engagement Requirements — Implications for Medicaid Enrollment & Adjudication

Recent policy changes have renewed focus on **Medicaid Community Engagement Requirements**, with significant implications for **eligibility, enrollment, renewals, and case adjudication**. While implementation will vary by state, the core operational shift is consistent: systems and workflows must now support **ongoing compliance tracking** in addition to traditional eligibility determinations.

Historically, Medicaid eligibility has been based on **income, household composition, residency, and categorical status**. **Community Engagement Requirements** add a new layer, requiring states to determine whether individuals must participate in **qualifying activities** (e.g., employment, job training, education, or volunteering), whether they **meet requirements or qualify for exemptions**, and what actions follow if they do not.

These requirements do not apply universally. **Exemptions**—often complex and time-limited—typically include individuals with **disabilities, caregivers, pregnant or postpartum individuals, older adults, individuals with certain medical conditions**, and others. As a result, eligibility becomes a **continuous adjudication process** rather than a one-time determination. With eligibility becoming a continuous process, systems will need to distinguish between routine compliance tracking and actions that may result in changes to coverage.

### Operational Impacts

Community engagement requirements shift Medicaid toward a compliance-tracking model, creating several challenges:

- **More frequent touchpoints:** States may need to assess compliance quarterly or more often, increasing case volume. This shift may significantly increase case volume and staffing demands, particularly during early implementation periods.
- **Expanded documentation burden:** New verification types (e.g., employment records, program participation, exemption proof) add complexity for both systems and caseworkers.
- **Increased risk of churn:** Eligible individuals may lose coverage due to reporting or documentation gaps, creating significant risk of procedural disenrollment if not proactively addressed.
- **Greater reliance on data integration:** Automated verification through workforce, benefits, and disability data sources will be critical to reduce burden and improve accuracy.

## Implications for AI

These changes are likely to increase operational strain on eligibility systems and staff. States may turn to **AI-enabled tools** to support **beneficiary engagement, document processing, exemption identification, risk flagging, and workflow prioritization**. When carefully governed, AI can help **reduce administrative burden** and **minimize procedural disenrollment**. However, its use must be paired with **strong oversight, transparency, and safeguards** to ensure accurate coverage decisions.

## What Does This Guide Include?

These use-case-specific best practice guides provide high level industry and consensus defined insights and recommendations for the application of responsible AI principles to a specific use case. The guide is organized by role (developer/implementer) and responsible AI principle areas where applicable (see figure below).

**This guide focuses on potential AI solutions in the space of Medicaid Adjudication of existing enrollees (Use Case 2).**

**A separate partner guide focuses on potential AI solutions in the space of Medicaid Enrollment of new beneficiaries (Use Case 1).**

The goal of developing and implementing AI in adjudication is to increase efficiency, reduce manual review workload, and maintain fairness by, among other means, keeping human oversight for complex or borderline cases. This includes AI aiding caseworkers supporting beneficiaries. AI can integrate data from multiple sources (employer records, volunteer logs, training program participation, community-based organizations) and flag cases for human review where information is incomplete or ambiguous.



### Useful, Usable & Effective

AI must solve specific problems, provide clear benefits, be easy to use, and perform reliably over time.



### Fairness & Bias Management

AI systems should treat individuals and groups consistently, minimizing unjustifiable differences in outcomes caused by issues in data, design, deployment, or use.



### Safe & Reliable

AI systems must not harm patients, requiring thorough testing, risk assessments, and continuous monitoring.



### Transparent & Accountable

Stakeholders must understand how an AI system works, its limitations, and who is responsible for its impact.



### Secure & Private

AI systems must protect patient data with strong security measures to prevent breaches and ensure confidentiality.

# Who Is This Guide For?

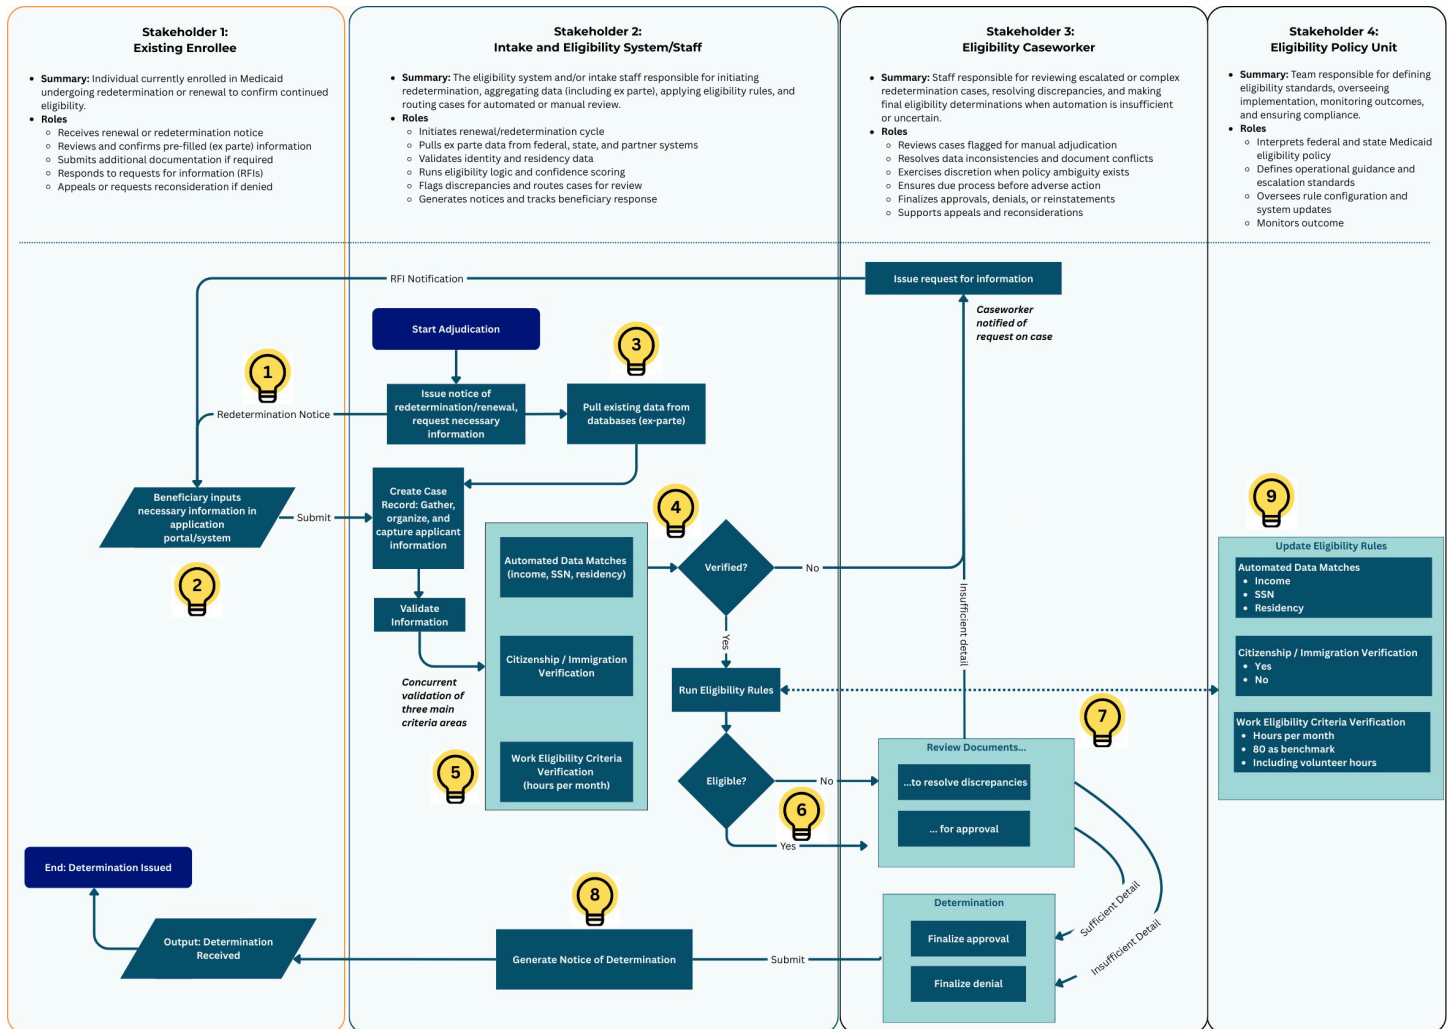


**Developers:** individuals involved in the software development process, including requirements gathering, design, coding, testing, and maintenance of software applications (derived from IEEE, 12207:2017)



**Implementers:** individual(s) responsible for the procurement, deployment, and/or overall realization of a system or component in accordance with a specified design (derived from IEEE 829 and IEEE 730)

## Workflow Example: Adjudication of Existing Enrollees



### Expanded View of Workflow



## Responsible AI Area 1)

### Usefulness, Usability, & Efficacy

#### Insights:

- **Seamless Integration into Workflows**

- AI tools are most effective when integrated into existing systems and workflows. Solutions that require extra steps, system switching, or do not account for state and county level controls are unlikely to be adopted or used consistently.

- **Human Oversight for High-Stakes Decisions**

- AI should support, not replace, difficult adjudication decisions. Human review of potential low-confidence denials and edge cases is critical before adverse actions to prevent errors, account for incomplete data, and maintain trust. Human review should include access to underlying information, the ability to override automated outputs, and clear documentation of decisions.

- **Accessible and Inclusive Adjudication Processes**

- Enrollees need multiple access options, including non-digital pathways. Literacy, language, and disability barriers should be addressed to ensure reliable use of AI-supported adjudication tools.

- **Targeted Support to Reduce Errors and Churn**

- AI adds significant value to beneficiaries and State adjudicators by improving documentation completeness, providing timely and relevant prompts, and helping prioritize actions. This reduces administrative errors and prevents unnecessary loss of coverage due to process failures.



#### Best Practices for Developers

1	<b>Design beneficiary-facing interfaces to be accessible on both mobile devices and computers, with low-bandwidth and offline-capable functionality where feasible.</b> <ul style="list-style-type: none"><li>● Maintain and clearly route to non-digital or low-tech options (phone/text, mail, in-person support) as a fallback when digital/AI-supported pathways fail.</li></ul>
2	<b>Embed clear checkpoints and user-controlled escalation mechanisms to route cases to human review when automation stalls or produces adverse eligibility outcomes (e.g., denial, termination, or loss of coverage).</b> <ul style="list-style-type: none"><li>● Prohibit fully automated eligibility denials and automated disenrollment in procurement and governance.</li><li>● Require documented human adjudication authority for final decisions and appeals.</li></ul>
3	<b>Support multilingual access aligned to state thresholds and accessibility standards (e.g., ADA, WCAG 2.2 AA, and Section 508), and incorporate low-literacy design principles.</b> <ul style="list-style-type: none"><li>● Pair AI-enabled language and accessibility features with explicit pathways to human assistance for eligibility support.</li></ul>

<p><b>4</b></p>	<p><b>Track beneficiary comprehension and response metrics (e.g., message open/response rates, successful document submission, time to resolution) and use results to iteratively improve communications and workflows.</b></p> <ul style="list-style-type: none"> <li>For improved comprehension and responsiveness, maintain consistent, standardized communications including tone and content, with allowances for configurable local elements (e.g., county contacts, terminology).</li> </ul>
<p><b>5</b></p>	<p><b>Integrate AI-supported adjudication capabilities directly into existing systems (e.g., eligibility platforms, EHRs, HIE-enabled tools) rather than requiring users to leave their primary workflow or learn parallel systems.</b></p> <ul style="list-style-type: none"> <li>Where full integration is delayed, unavailable, or introduces operational risk, use interim workflows that minimize disruption and duplicative data entry while progressing toward deeper integration over time.</li> <li>When integrating EHR/HIE data, clinical information should be used to identify and protect likely exemptions or prompt follow-up, not as a standalone basis for adverse eligibility determinations, especially when records are incomplete, ambiguous, or not aligned with administrative documentation standards.</li> </ul> <p style="text-align: right;"><i>*See Appendix 1 for a useful example</i></p>
<p><b>6</b></p>	<p><b>Use targeted, context-aware alerts that surface only when action is required, allowing users to override or proceed with informed acknowledgment.</b></p> <ul style="list-style-type: none"> <li>Use event- and deadline-driven alerts that prioritize high-urgency cases and tailor outreach based on beneficiary context (where available and permitted), while avoiding excessive or confusing notifications.</li> <li>Implement AI-driven documentation completeness checks and proactive beneficiary reminders of upcoming redeterminations tailored to what is missing and why it matters.</li> </ul> <p style="text-align: right;"><i>*See Appendix 1 for a useful example</i></p>
<p><b>7</b></p>	<p><b>Implement triage and prioritization mechanisms (queues, severity scoring, and workload balancing) so AI outputs are actionable at realistic staffing levels.</b></p>



## Best Practices for Implementers

<p><b>A</b></p>	<p><b>Simplify access across systems by minimizing logins, navigation steps, and credential requirements.</b></p> <ul style="list-style-type: none"> <li>Where single sign-on is infeasible, prioritize workflow integration and reduced administrative burden for staff.</li> <li>Require vendors to integrate AI-supported functions into existing eligibility and caseworker workflows and validate that integration is workable in the real state/county environment.</li> </ul>
-----------------	--

<p><b>B</b></p>	<p><b>Offer multiple support pathways—including navigators and caseworkers—to assist beneficiaries when AI-supported tools are inaccessible, confusing, or incomplete.</b></p> <ul style="list-style-type: none"> <li>● Train staff to recognize AI escalation triggers, particularly for denials or eligibility loss, and ensure rapid-response human review pathways.</li> <li>● Establish a human-in-the-loop requirement for adverse actions (e.g., disenrollment) with defined review steps and grace periods when data is missing, lagged, or conflicting. This review should include access to underlying information, authority to override automated outputs, and documentation of decision-making.</li> <li>● Clearly define and surface AI escalation triggers and system limitations so implementers can build reliable human review pathways based on known risk conditions.</li> </ul> <p style="text-align: right;"><i>*See Appendix 1 for a useful example</i></p>
<p><b>C</b></p>	<p><b>Where county-level coordination and eligibility determination exists, balance statewide consistency with county-level customization in outreach, timelines, and points of contact.</b></p> <ul style="list-style-type: none"> <li>● Define and document state/county operational interpretations and edge-case handling before rollout (including who decides, how updates are issued, and how caseworkers are supported) to reduce inconsistent determinations.</li> </ul>
<p><b>D</b></p>	<p><b>AI-supported workflows should proactively be focused on prevention (early outreach, document reminders, and rapid correction pathways) ahead of a determination.</b></p> <ul style="list-style-type: none"> <li>● Define churn-reduction objectives (e.g., reducing procedural disenrollment) as explicit program success metrics and track outcomes.</li> </ul>



## Responsible AI Area 2)

### Fairness & Bias Management

#### Insights:

- **Protecting Groups in High-Risk Decisions**
  - Exemption and work-related eligibility decisions carry high fairness risk. AI should help identify likely exemptions and protect individuals at higher risk, not just enforce rules.
- **Governance of Models and Rules**
  - Model and rule choices directly affect fairness. Systems should avoid black-box decisions and ensure logic stays valid as policies, data, and populations change.
- **Fairness Beyond Model Performance**
  - Fairness depends not only on AI system accuracy, but also on access to documentation and communication channels, which can create unequal outcomes across groups.



## Best Practices for Developers

1

**Support automated identification of likely exemptions (e.g., medically frail or caregiver status) using transparent logic, audit trails, and optional human confirmation.**

- Prioritize models/rules that surface likely exemptions and risk-of-harm signals (e.g., disability, caregiver status, postpartum) and require human confirmation before adverse determinations when exemption evidence is low-confidence or ambiguous.
- Where advanced analytics or Natural Language Processing are used to identify exemptions or risk signals, distinguish probabilistic signals from deterministic, and when signals for exemption are of a high-confidence or clear, auto-route for streamlined approval.

*\*See Appendix 1 for a useful example*



## Best Practices for Implementers

A

**Establish governance processes (in coordination with developers) to review, approve, and periodically reassess which AI models, algorithms, or rule sets are used for adjudication.**

- Require transparency in how fairness and bias are defined and assessed for relevant populations.

*\*See Appendix 1 for a useful example*

B

**Monitor disparities in response rates, documentation completion, and appeal success across demographic groups and adjust outreach and support to close gaps.**



## Responsible AI Area 3)

### Safety & Reliability

#### Insights:

- **Surfacing Uncertainty and Preventing Harm**
  - Safety depends on identifying missing, delayed, or conflicting data. Systems should surface uncertainty and avoid defaulting to denial when information is incomplete.
- **Flexible Systems for Policy Change**
  - Policies and interpretations vary across states and over time. Systems should be adaptable and avoid hard-coded assumptions that can lead to errors.
- **Safe Rollout and Workforce Readiness**
  - AI should be introduced gradually with strong training, clear escalation pathways, and alignment across state and county operations to reduce risk.

- **Operational Resilience and Capacity**

- AI can increase demand on support systems. Rollouts should account for call center capacity and ensure continuity of service during high-volume periods.



## Best Practices for Developers

<b>1</b>	<b>Provide dashboards tracking error rates, appeals, disenrollments, and anomalies to support oversight and improvement.</b> <ul style="list-style-type: none"><li>● Detect and surface data uncertainty (e.g., missing records, 30-90 day lags, conflicting sources) and prevent automated adverse actions when uncertainty exceeds predefined thresholds. When key data is missing, outdated, or conflicting, systems should prioritize clarification or human review.</li></ul> <p style="text-align: right;"><i>* See Appendix 1 for a useful example</i></p>
<b>2</b>	<b>Use modular, configurable architectures adaptable to state and county rules without full system redesign.</b> <ul style="list-style-type: none"><li>● Build configurable rule and workflow components that can be updated quickly as policy guidance changes, including clear documentation of where interpretation is required and how updates are governed.</li></ul>
<b>3</b>	<b>Define safe-default behaviors (e.g., pause, request clarification, route to human review) when data is missing or contradictory, and prohibit default-to-denial logic in these scenarios.</b>



## Best Practices for Implementers

<b>A</b>	<b>Establish a standardized intake and evaluation process for AI vendor solutions (including minimum safety, transparency, and auditability requirements) for timely and successful adoptions.</b> <ul style="list-style-type: none"><li>● Where AI implementations are found successful, participate in inter-state sharing of best practices to replicate proven, safe models.</li></ul>
<b>B</b>	<b>Embed navigators and outreach roles throughout AI-supported adjudication to maintain transparency and beneficiary engagement.</b> <ul style="list-style-type: none"><li>● Define navigator/caseworker escalation criteria (e.g., low-confidence matches, high-urgency needs, pending adverse action).</li></ul>

**C****Adopt phased or pilot-based deployment strategies for AI-supported adjudication, with clearly defined success metrics and rollback mechanisms before scaling statewide.**

- Where county-level determination exists, establish a coordinated state-county approach and governance process including vendor oversight and assessing of workforce readiness.
- Require safeguards that block denials during active appeals, avoid penalizing short-term compliance gaps, and enable rapid reinstatement when errors are identified.
- Provide mechanisms for staff to report identity matching failures and route those reports into corrective workflows.

*\*See Appendix 1 for a useful example*

**D****Conduct operational readiness and stress-testing (call center volume, navigator staffing, language coverage) before scaling AI-supported adjudication workflows.**

- Implementation planning should account for call center capacity, staffing, and support resources during high-volume periods.



## Responsible AI Area 4)

### Transparency

#### Insights:

- **Clear Communication and User Understanding**
  - Beneficiaries need clear, simple explanations of how decisions are made, what actions are required, and how to fix issues without navigating complex processes.
- **Accessible Appeals and Due Process**
  - Beneficiary trust depends on clear appeal pathways, understandable exemption processes, and timely correction when errors occur.
- **Visibility into Data and System Behavior**
  - Systems should make data sources, limitations, and uncertainty visible, especially when data is incomplete, outdated, or conflicting.
- **Defined Accountability and Oversight**
  - Clear responsibility should be defined across agencies, counties, vendors, and intermediaries, supported by audit trails, performance monitoring, and transparent reporting.



## Best Practices for Developers

<p>1</p>	<p><b>Enable two-way communication, real-time case status updates, and alerts when information is missing or under review.</b></p> <ul style="list-style-type: none"><li>● Ensure two-way communication supports multilingual use and reaches both beneficiaries and staff as appropriate.</li></ul> <p><i>* See Appendix 1 for a useful example</i></p>
<p>2</p>	<p><b>Provide plain-language explainability and visible data-use indicators showing where and how automation influences determinations, especially for complex cases.</b></p> <ul style="list-style-type: none"><li>● Provide aggregate dashboards and exportable reporting at key forks (requests for info, denials, exemptions, reinstatements) with demographic and geographic breakdowns appropriate for oversight and monitoring.</li><li>● Make this information readily available (e.g., via a short brochure or explainer) and accessible upon request, rather than requiring detailed AI explanations during every beneficiary interaction.</li></ul> <p><i>* See Appendix 1 for a useful example</i></p>
<p>3</p>	<p><b>Design systems to identify and clearly surface data gaps, quality issues, timeliness limitations, and data-sharing restrictions that could affect eligibility or exemption determinations.</b></p> <ul style="list-style-type: none"><li>● Surface source provenance and uncertainty indicators (timeliness, conflicts, missing records) directly in caseworker views and audit logs so staff understand when automation is reliable and when it is not.</li></ul>
<p>4</p>	<p><b>Provide beneficiary-facing status views that clearly show what is needed, why it is needed, and how to resolve it, including confirmation when documents are received and processed.</b></p>



## Best Practices for Implementers

<p>A</p>	<p><b>Maintain clearly communicated appeal and exemption pathways with defined timelines and human support.</b></p> <ul style="list-style-type: none"><li>● Provide plain-language explanations of data sources, sharing practices, and AI use within adjudication workflows.</li><li>● Ensure notices and appeal materials clearly explain data sources, what is missing or inconsistent, and the human rationale for adverse actions, using plain language and a clear correction/appeal pathway with timelines.</li><li>● Ensure early-stage human intervention does not create bottlenecks by enabling AI-supported reconsideration opportunities to be routed efficiently to caseworkers for reevaluation before formal appeals.</li></ul>
----------	---

<p><b>B</b></p>	<p><b>Perform regular audits and share aggregated findings to verify safety, reliability, and fairness.</b></p> <ul style="list-style-type: none"> <li>● Maintain complete audit trails (inputs, model/rule version, outputs, overrides, and human decisions) and provide audit access mechanisms suitable for oversight, compliance, and litigation readiness.</li> <li>● Where feasible, provide actionable error findings to relevant provider organizations and operational stakeholders to support corrective action.</li> </ul>
<p><b>C</b></p>	<p><b>Clearly define roles and responsibilities for data stewards (e.g., HIEs, data utilities) related to data quality, governance, consent management, and audit support within AI-assisted adjudication workflows.</b></p> <ul style="list-style-type: none"> <li>● Assign named owners for AI governance (including data stewardship, policy interpretation, and model/rule oversight) and define decision rights for updates, incident response, and stakeholder communication.</li> </ul>
<p><b>D</b></p>	<p><b>Establish documented procedures for error remediation, beneficiary notification, and corrective action when AI-supported adjudication contributes to incorrect or harmful outcomes.</b></p> <ul style="list-style-type: none"> <li>● Implement documented error remediation procedures that include corrective action tracking and rapid reinstatement workflows when AI-supported processes contribute to harm or incorrect outcomes.</li> </ul>
<p><b>E</b></p>	<p><b>Establish and publish a responsibility matrix (e.g. RACI) covering decision-making, monitoring, incident response, and remediation for all AI-supported adjudication components.</b></p> <p style="text-align: right;"><i>*See Appendix 1 for a useful example</i></p>



## Responsible AI Area 5)

### Security & Privacy

#### Insights:

- **Purpose-Limited Data Use and Sharing**
  - Data collection and sharing should be limited to what is necessary, with clear controls to reduce privacy risk as data is aggregated across systems.
- **Cross-Agency Data Readiness and Governance**
  - Effective deployment depends on coordinated data governance across agencies, including managing consent, access, and data quality challenges.
- **Secure Identity and Data Handling**
  - Identity matching and data reconciliation are high-risk areas. Strong safeguards are needed to prevent incorrect eligibility actions or unauthorized access.
- **Security Against Emerging Threats**
  - AI systems increase exposure to fraud and manipulation. Security should include protections against identity fraud, adversarial inputs, and misuse of sensitive eligibility data.



## Best Practices for Developers & Implementers

<b>A</b>	<p><b>Automate retrieval of high-quality, authoritative data (e.g., ex parte verification via federal, state, and HIE sources) with transparency around provenance.</b></p> <ul style="list-style-type: none"><li>● Pre-fill inputs while allowing beneficiaries and staff to review, correct, and confirm information prior to submission.</li><li>● Set minimum security and privacy requirements for cross-agency data access (least privilege, auditing, retention limits) and ensure beneficiaries understand what data is used and why.</li></ul> <p style="text-align: right;"><i>*See Appendix 1 for a useful example</i></p>
<b>B</b>	<p><b>Conduct data readiness assessments across participating entities (e.g., state agencies, counties, and relevant partners) to identify gaps, duplication, and consent barriers.</b></p> <ul style="list-style-type: none"><li>● Include consent/access constraints and operational feasibility (who can access what data, under what authority, and how gaps will be handled without adverse action).</li></ul>
<b>C</b>	<p><b>Implement safeguards for patient and beneficiary matching, including confidence scoring, duplicate detection, and flagging of low-confidence matches before automated determinations are made.</b></p> <ul style="list-style-type: none"><li>● Implement robust identity verification and reconciliation controls and log access/changes to support investigation and accountability.</li></ul> <p style="text-align: right;"><i>*See Appendix 1 for a useful example</i></p>
<b>D</b>	<p><b>Conduct threat modeling for identity fraud and adversarial submissions (fake documents, synthetic identities) and implement detection and escalation controls appropriate for Medicaid workflows.</b></p> <p style="text-align: right;"><i>*See Appendix 1 for a useful example</i></p>
<b>E</b>	<p><b>Establish privacy governance that limits use and re-sharing of sensitive eligibility signals to the minimum necessary, with clear retention policies and auditability.</b></p>

## AI Ideas

Given the best practices identified and the additional context of the Adjudication Use Case example workflow, here are several community developed ideas for where AI may benefit the process.



## Intelligent Request for Clarifying Information or Notice Simplification

- **Where:** "Issue notice of redetermination / request necessary information"
- **What AI does:**
  - Natural language simplification of notices
  - Personalization of required documentation based on known profile
  - Translation & readability optimization (CLAS-aligned)
- **Why it matters:**
  - Higher first-pass completeness
  - Fewer requests
  - Reduced churn



## Document Intake & Classification

- **Where:** "Beneficiary inputs necessary information"
- **What AI does:**
  - OCR + document type classification
  - Real-time completeness checks ("You uploaded pay stub but not employer letter")
  - Fraud anomaly detection (basic level)
- **Why it matters:**
  - Fewer back-and-forth cycles
  - Reduced caseworker manual triage
  - Lower improper denial risk



## Automated Data Aggregation (Ex-Parte Expansion)

- **Where:** "Pull existing data from databases (ex-parte)"
- **What AI does:**
  - Intelligent matching across income, residency, SNAP, UI data
  - Confidence scoring for data reliability
  - Real-time gap detection ("Income source outdated by 75 days")
- **Why it matters:**
  - Reduced manual verification
  - Faster verified cases
  - Safer uncertainty routing



## Automated Data Matches (Income / SSN / Residency)

- **Where:** "Automated Data Matches"
- **What AI does:**
  - Predictive anomaly detection
  - Risk scoring for mismatches
  - Differentiating data error vs eligibility failure
- **Why it matters:**
  - Fewer unnecessary caseworker escalations
  - Reduced false denials



## Community Engagement Requirement Criteria Verification

- **Where:** Community Engagement Requirement Logic
- **What AI does:**
  - Pattern detection across wage records
  - Flagging likely exemption (caregiver, disability)
  - Identifying gig-economy irregular income

- **Why it matters:** Surface likely exemptions before determination.



### Confidence-Based Routing Model

- **Where:** At point of document review and resolution of discrepancies
- **What AI does:** Establish eligibility confidence tiers:
  - $\geq 95\%$  → streamlined approval
  - 80–94% → human confirmation
  - $< 80\%$  → unclear, more information needed
- **Why it matters:**
  - Speeds low-risk approvals
  - Protects high-risk cases



### Smart Triage Queue

- **Where:** “Review Documents...”
- **What AI does:**
  - Risk-based case prioritization
  - Urgency scoring (pregnancy, surgery)
  - Discrepancy clustering
- **Why it matters:**
  - Reduced backlog
  - Better allocation of human review



### Determination Drafting Assistant

- **Where:** “Generate Notice of Determination”
- **What AI does:**
  - Draft plain-language determinations
  - Ensure required regulatory language included
  - Check internal consistency
- **Why it matters:** Timely notice of pending determination result to allow beneficiary corrections.
  - Reduced drafting time
  - Fewer appeal reversals



### Rule Simulation Engine

- **Where:** “Update Eligibility Rules”
- **What AI does:**
  - Simulate impact of rule changes before deployment
  - Detect disparate impact
  - Predict churn or workload increase
- **Why it matters:**
  - Prevent policy-induced operational spikes
  - Advance fairness testing

# Appendix 1: Examples

A list of practical examples which elaborate on specific terminology and concepts found in best practice statements. Numbers and letters refer to the Best Practice they illustrate from the corresponding principle area.

<p>5</p>	<p><b>Human-Centered Workflow Design</b></p> <p><i>A human-centered workflow:</i></p> <ul style="list-style-type: none"><li>● Fits into existing caseworker screens and daily processes,</li><li>● Reduces duplicate data entry,</li><li>● Surfaces only actionable alerts.</li></ul> <p><i>Poor design: A separate AI portal requiring separate login.</i></p> <p><i>Good design: AI insights embedded in the existing eligibility system screen.</i></p>
<p>6</p>	<p><b>Avoiding Alert Fatigue</b></p> <p><i>If every case generates alerts, staff will override them reflexively.</i></p> <p><i>Good practice:</i></p> <ul style="list-style-type: none"><li>● Alert only when action is required.</li><li>● Suppress repeat alerts.</li><li>● Allow meaningful override with documentation.</li></ul>
<p>1</p>	<p><b>Fairness Testing in Medicaid AI</b></p> <p><i>Fairness testing evaluates whether model performance differs across groups.</i></p> <p><i>Common metrics:</i></p> <ul style="list-style-type: none"><li>● Error rate parity (Are denials disproportionately incorrect for certain groups?)</li><li>● Calibration (Is risk scoring equally accurate across subgroups?)</li><li>● False negative rates (Are certain populations under-identified for exemptions?)</li></ul> <p><i>Fairness testing should occur:</i></p> <ul style="list-style-type: none"><li>● Before deployment,</li><li>● After major model updates,</li><li>● On a recurring basis.</li></ul>
<p>1</p>	<p><b>Confidence thresholds</b></p> <p><i>AI systems should not treat all outputs equally, for example:</i></p> <ul style="list-style-type: none"><li>● If identity match confidence <math>\geq 95\%</math> → proceed</li><li>● 80–94% → human review</li><li>● &lt; 80% → unclear, more information needed</li></ul> <p><i>This prevents low-confidence automation from triggering coverage loss.</i></p>

1

## What should be reported publicly?

Consider:

- Number of AI-assisted determinations,
- Denial rates,
- Exemption identification rates,
- Appeal rates,
- Breakdown by geography and key demographic indicators.

*Transparency builds trust and enables oversight.*

2

## Beneficiary-facing explainability

*Beneficiaries do not need model math. They need:*

- What data was used,
- What is missing,
- What action they can take,
- Who to contact.

*Explainability should empower correction, not overwhelm.*

B

## What does human-in-the-loop mean in practice?

*Human-in-the-loop means more than glancing at an output. It means:*

- AI may recommend or flag a case,
- But a trained caseworker makes the final determination before any adverse action,
- Especially where data is missing, contradictory, or high-risk.

*For example if wage data is missing or 60 days old, the system flags the case as "Data Uncertain - Review Required" and prevents automatic disenrollment until a human reviews the file.*

A

## What does "no black box" mean?

*A system is a black box if:*

- Decision logic cannot be explained,
- Variables cannot be audited,
- Model versions are not tracked.

*Minimum expectations:*

- Version control,
- Documented decision logic,
- Audit logs,
- Ability to explain why a determination was made.



## Safe defaults

*When data is missing or conflicting, systems should default to:*

- *Pause,*
- *Request clarification,*
- *Route to human review.*

*Never default to denial.*



## What is a RACI matrix?

*RACI =*

- *Responsible (does the work)*
- *Accountable (owns the outcome)*
- *Consulted (provides input)*
- *Informed (kept updated)*

*For AI adjudication:*

- *Vendor may be Responsible for model updates,*
- *State agency Accountable for determinations,*
- *Counties Consulted,*
- *CMS Informed.*

*Clear ownership prevents governance gaps.*



## Zero Trust in Medicaid AI

*Zero Trust means:*

- *No user or system is trusted by default,*
- *Access is continuously verified,*
- *Permissions are limited to what is strictly necessary (“least privilege”).*

*Example:*

*A vendor system should not access full beneficiary records if only eligibility status is needed.*



## Identity matching safeguards

*Good identity controls include:*

- *Confidence scoring,*
- *Duplicate detection,*
- *Manual review triggers for low-confidence matches,*
- *Logging of all identity overrides.*

*Incorrect identity matches can lead to wrongful termination or data exposure.*



## Threat modeling

*Threat modeling asks:*

- *How could someone game this system?*
- *Could fake paystubs bypass controls?*
- *Could automated bots flood document uploads?*

*Security planning should include:*

- *Fraud detection,*
- *Anomaly monitoring*
- *Escalation pathways*

## Appendix 2: Methods

Best practice statements are collected from work group presentations and discussions. To ensure alignment across stakeholders, CHAI uses a multi-phase consensus process for Best Practice Statements (BPS) generated through work group activities:

### Phase 0: Convening Tiger Team + Presentations with Discussions (August - October 2025)

- **Purpose:** Define Use Cases and garner expertise from guest presenters/discussions.
- **Format:** Regular biweekly meetings, alternating by use case

### Phase 1: Initial Consensus Check (November 2025)

- **Purpose:** Initial draft of BPS and gauge initial agreement.
- **Voting Options:** *Include / Include Contextually / Exclude / Abstain*
  - *Obtained through asynchronous surveying of extended tiger team (reviewing responsibilities), and live review during core tiger team meetings (drafting responsibilities).*
- **Decision Rules:**
  - Simplified consensus model.
    - BPS with any Exclude votes were prioritized for review during core tiger team meetings for feedback.
      - If  $\geq 2/3$  vote "Exclude" → Automatically excluded
    - Survey responses for additional context were also reviewed and incorporated.

### Phase 2: Revote with Revisions (December 2025)

- **Purpose:** After redrafting, re-evaluate updated BPS.
  - Includes enhancements based on recent presentations and discussions.
- **Voting Options:** *Include / Exclude / Abstain*, with an optional comment field.
  - *Obtained through asynchronous surveying of extended tiger team (reviewing responsibilities), and live review during core tiger team meetings (drafting responsibilities).*
- **Decision Rules:** Same simplified consensus model applied.
- **Outcome:** Second round of feedback achieved, adding greater detail to statements.
  - Initial drafts of use case workflows were also created as a result of discussions

### Phase 3: Live Discussion and Finalizing (January-April 2026)

- **Purpose:** After redrafting, final round of feedback on BPS.
  - Includes enhancements based on recent presentations and discussions.
- **Voting Options:** *Include / Exclude / Abstain*, with an optional comment field.
  - *Obtained through asynchronous surveying of extended tiger team (reviewing responsibilities), and live review during core tiger team meetings (drafting responsibilities).*
- **Decision Rules:** Same simplified consensus model applied.
  - Facilitated group discussion of flagged BPS.
- **Outcome:** Final BPS made based on discussion and final review results.

## Appendix 3: Thank You and Contributors

We want to start by thanking every individual who showed interest, participated, listened, and came along with us in the early stages of our work. CHAI is, at its core, a convener and a member-driven non-profit. We are so grateful to be on this journey with you towards responsible AI in health for all. Your experiences, your feedback, your contributions, all make us who we are and help bring us to where we need to be.

For those who want to be credited directly by name, please reach out to us at [greg@chai.org](mailto:greg@chai.org) to request contribution credit for the Tiger Team on AI & Medicaid. Below is a list of organizations who had at least one individual who showed interest and/or participated in the tiger team.

If you want to learn more about our work groups (current and future), or have feedback on this effort, products, or services, please contact our program management team via Greg Shemancik ([greg@chai.org](mailto:greg@chai.org)).

### Participating Organizations (Extended, Core, and Co-Chairs for the Tiger Team)

- Asian Health Services
- BridgeHealthAI
- Centene [Co-Chair]
- Community Health Network
- Community–University Health Care Center
- CompleteCare
- Cornell Scott Hill Health Center
- CRISP Shared Services
- Delaware Valley Community Health
- DPE International
- Duke–Margolis Institute for Health Policy
- Escher Health
- HEAL, Stanford University
- HealthTech 4 Medicaid (HT4M) [Co-Chair]
- Innovation Hub Advisors
- Iowa Primary Care Association (Iowa PCA)
- ITO Health
- Logick.io
- MettaHealth Partners
- Mirza
- MPCA
- National Association of Community Health Centers (NACHC) [Co-Chair]
- National Council for Mental Wellbeing
- National Health Council

- Nextiva
- Ochsner Health
- Pair Team [Co-Chair]
- Prime Health
- Reya Health
- SC Primary Health Care Association
- SciEncephalon Corp
- Society of General Internal Medicine
- Synergist Technology
- UiPath
- Valley Health Partners Community Health Center