International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 02 | Feb - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

# Modernizing Medicaid Rules with Oracle Intelligent Advisor (OIA) Cloud

**Malini Balasubramanyam**

**Program:** Rules movement to Cloud as part of a Technical Upgrade Release

**Timeline:** November 2024 – May 2025

## Executive Summary

Over seven months, AHSS executed a strategic modernization of a large, legacy rules estate—over **10,000 complex Medicaid rules**—transitioning from an embedded, on-premises rules engine approaching end-of-life to **Oracle Intelligent Advisor (OIA) SaaS**. The initiative emphasized deep mastery of OIA's rule design and orchestration model and delivered a cloud-ready rules architecture that eliminates inter-budget dependencies, resolves evaluation loops, and preserves decision parity for actively used budgets. The result is a scalable, maintainable, and auditable rules corpus that aligns with cloud constraints while retaining functional fidelity where it matters to clients.

## Background & Motivation

- **Retire end-of-life embedded Rule logic** and reduce operational risk by moving to OIA SaaS based Rule logic which will help to leverage **cloud-native rule governance, transparency, and auditability.**
- **Standardize Medicaid budgeting determinations** across jurisdictions and households.
- The goal was to keep outcomes the same where clients actively rely on them, while cleaning up how the rules are designed.

## Scope & Objectives

- **Primary scope:** Medicaid rulebase migration to OIA SaaS, including all active budgets and household/person scoping.
- **Objective 1:** Remove rule loops and implicit budget dependencies which was allowed in embedded Rules.
- **Objective 2:** Infer business-critical inputs inside rules based on explicit business conditions.
- **Objective 3:** Align outputs between embedded and cloud rules for all actively used business conditions / scenarios.
- **Objective 4:** Establish durable rule design patterns (decision trees, ranking/prioritization, scoping) for future maintenance.

## Key Technical Challenges & How We Solved Them

- **Medically Needy property budgeting:** Some conditions were resolving as "unknown," creating loops across budgets. We refactored so each budget is evaluated individually and resolves to a clear Pass/Fail.
- **Embedded vs. Cloud behavior:** The old system allowed one budget to infer another. In the cloud, we removed those dependencies and used explicit evaluation order and ranking instead.
- **Data handling:** We stopped front-loading data. The rules now infer what they need based on business logic and scope that across household members.
- **People relationships:** We now prove and store all relevant relationships once within the rulebase based on identifiers and conditions.

- **Reference table updates:** The Budget Category Reference Table (Excel) was updated to reflect the new inference approach.
- **Result parity controls:** After removing loops, a few attributes diverged between embedded and cloud. We corrected those by tightening the decision trees and adding controlling logic, so that results align.
- **Full Scope Coverage dependency:** Removed. Business Scenarios now function independently and are resolved via ranking / priority ordering.

---

## What to Expect Operationally

- With dependencies removed, **more budgets may be created** than in the embedded system. That's expected. Final selection is now driven by ranking and explicit criteria.
- independently and are resolved via ranking / priority ordering.

---

## Solution Architecture

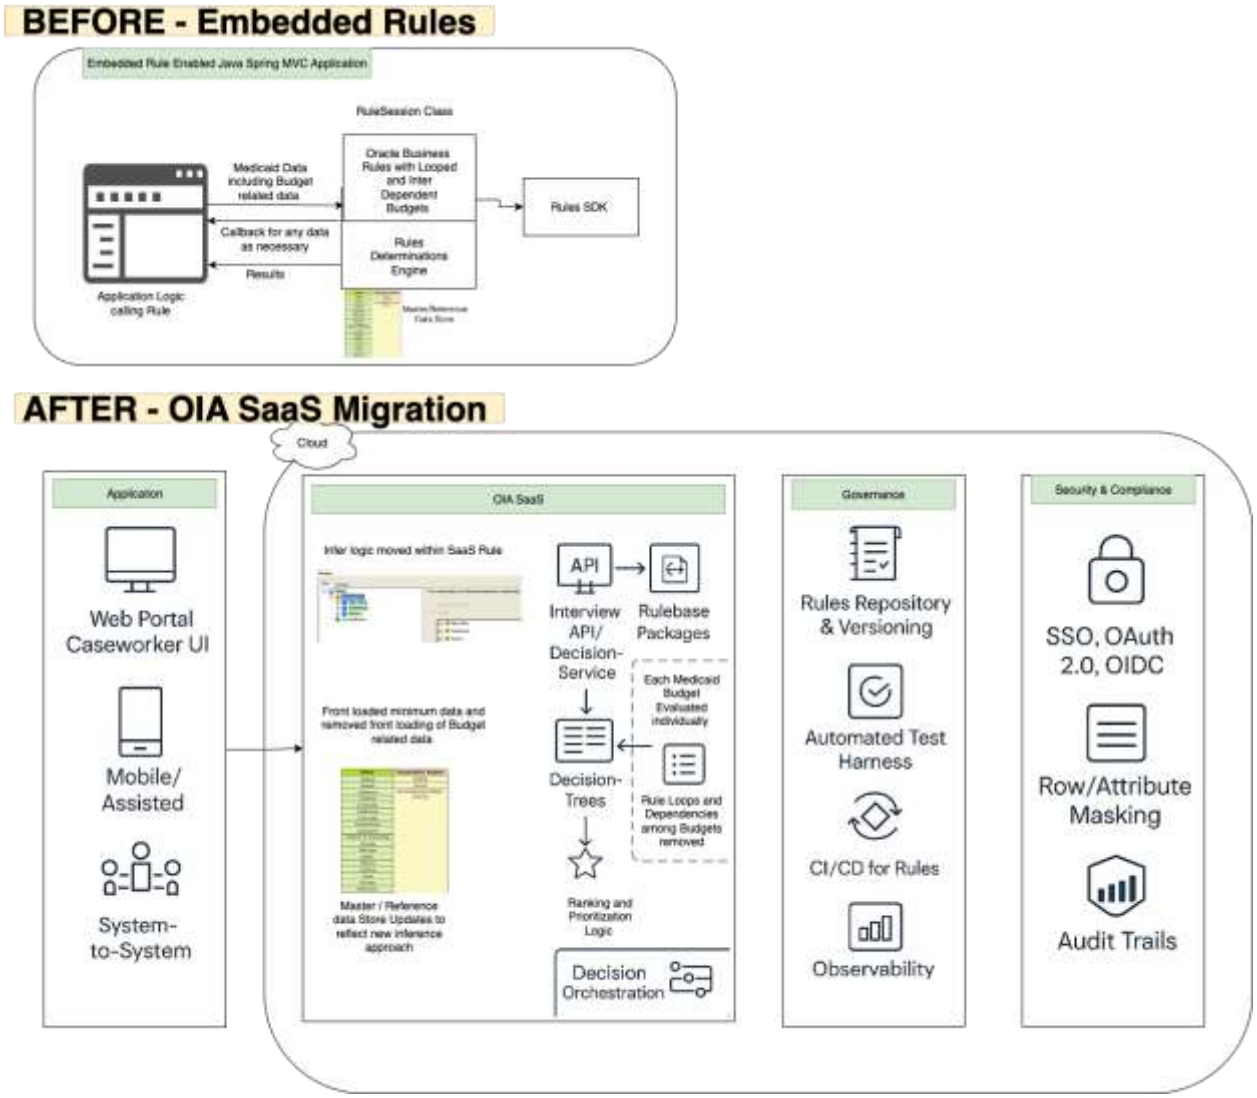


---

## Metrics (Effort and Outcomes)

- **Rules transitioned:** 10,000+ (portfolio scale; highly complex).
- **Budgets covered:** 54 total; **46 validated to parity** with embedded.

- **Parity rate (validated budgets): ~85%** (46/54).
- **Timeline:** ~7 months (Nov 2024–May 2025).
- **One-time data alignment:** Reference Table data (master data) updated to match new rule inference.
- **Loop elimination:** All cross-budget loops removed for the 46 validated budgets.
- **Decision tree tuning:** Applied where minor attribute differences were observed to regain parity.

---

## Governance and Quality

- Rules were aligned to use cases (e.g., medically needy budgeting) during reviews.
- Budget **hierarchy and priority** were validated to ensure predictable outcomes.
- Removed budget dependencies and replaced them with deterministic sequencing and ranking.
- Decision trees serve as the main traceable artifact for audits and targeted corrections.

---

## Lessons Learned

- Design rules to be **independent** from the start; don't rely on one budget to determine another.
- When removing dependencies, expect **more generated budgets** and handle selection via ranking.
- Keep **reference data** synchronized with rule logic; small mismatches lead to noisy results.
- **Decision trees** are effective for pinpoint parity corrections and explainability.

---

## Next Steps

1. Re-confirm the **priority ordering** against the latest use case documents.
2. Publish a **sustainment guide** (patterns for scoping, inference, decision trees, and ranking).
3. Keep monitoring for any residual attribute drift and fix via small, targeted rule updates.