International Journal on Science and Technology (IJSAT)
E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

AI-Augmented Agentforce Architectures for Predictive Healthcare Claims Adjudication in Salesforce Data Cloud

Susil Sahu

Independent Enterprise Salesforce Architect – Healthcare & Insurance Domain

Abstract

Healthcare payer organizations are under increasing pressure to improve claims turnaround time, reduce administrative burden, and maintain strong compliance controls. This paper presents a practitioner-oriented architecture for predictive healthcare claims adjudication using Salesforce Data Cloud and Agentforce in a governed, human-in-the-loop operating model. The proposed approach unifies claims, member, provider, and service data to improve contextual visibility, identify likely exception-prone claims earlier, and support adjudicators with workflow-aware recommendations. Rather than replacing adjudication authority, the model uses AI to strengthen triage, prioritization, and case understanding while preserving auditability and human review. The paper outlines the architectural components, workflow design, and governance measures needed to operationalize this model in a healthcare payer environment.

Keywords

Salesforce Data Cloud; Agentforce; Healthcare Claims Adjudication; Human-in-the-Loop AI; Enterprise Architecture; Healthcare Payer Modernization

1. Introduction

Claims adjudication remains one of the most operationally significant functions in healthcare payer organizations. It affects payment speed, provider relationships, member experience, and administrative cost. Yet many payer environments still rely on fragmented data, disconnected workflows, and limited contextual visibility during exception handling.

This paper proposes a practical architecture in which Salesforce Data Cloud acts as a unification layer for claims-related context while Agentforce supports prediction, triage, and decision assistance. The goal is not full automation. The goal is a stronger, more responsive, human-in-the-loop adjudication model.

2. Problem Context

A substantial share of claims can be processed through deterministic rules when data is complete and consistent. The operational challenge grows when claims involve missing information, prior

authorization dependencies, conflicting eligibility details, duplicate submissions, or unusual financial exposure. In these scenarios, adjudicators often spend more time assembling context than making decisions.

This creates a strong need for systems that surface relevant information earlier, prioritize likely exception-heavy cases, and support faster, more consistent adjudication without weakening compliance controls.

3. Proposed Architecture

The proposed architecture consists of four main layers: data ingestion and harmonization, context unification, AI-assisted recommendation, and governed workflow orchestration. Salesforce Data Cloud aggregates operational signals from claims, member, provider, and service environments. Agentforce then uses this curated context to support exception prediction, summarization, and next-step guidance for adjudication teams.

A key design principle is bounded assistance. AI-generated guidance may support triage and review, but final claim decisions remain under human authority. This preserves both operational trust and regulatory defensibility.

4. Architecture Summary Table

- Data Layer: Claims, member, provider, and service records are ingested and harmonized for contextual use.
- Context Layer: Relevant case context is assembled for reviewer visibility and prioritization.
- AI Assistance Layer: Agentforce supports exception identification, summaries, and workflow recommendations.
- Governance Layer: Human review, access controls, explainability, and audit logging remain mandatory.

5. Workflow and Governance

The workflow begins when a claim enters the operating environment. Relevant records are assembled, screened through rule-based logic, and enriched with contextual indicators before AI-generated assistance is introduced. This may include exception likelihood, recommended routing, and a concise explanation of why the claim requires closer attention.

Governance controls include role-based access, controlled context retrieval, explicit review boundaries, evidence capture, and audit logging. These measures allow AI assistance to improve productivity while preserving a clear accountability structure.

6. Strategic Value

This model can reduce manual search effort, improve queue prioritization, and support more consistent handling of exception-heavy claims. It also creates a stronger operating foundation for payer modernization by connecting claims processing to a broader, more contextual service architecture.

From a professional standpoint, the topic reflects the intersection of Salesforce architecture, healthcare operations, and governed enterprise AI. That makes it highly relevant to current payer transformation initiatives.

7. Conclusion

Predictive claims adjudication should not be framed as uncontrolled automation. In healthcare payer environments, the more credible approach is AI-augmented adjudication built on strong data unification, workflow discipline, and human oversight. Salesforce Data Cloud and Agentforce together create a practical foundation for this model when deployed with appropriate governance and operational boundaries.

References

1. Salesforce, "Data Cloud Overview," Salesforce Documentation.
2. Salesforce, "Agentforce Overview and Prompt-Based Assistance," Salesforce Documentation.
3. Salesforce, "Health Cloud Overview," Salesforce Documentation.
4. Salesforce, "Security and Identity for Salesforce Platform," Salesforce Documentation.
5. JMIR Medical Informatics, public materials on healthcare digital workflow and informatics practice.
6. AMIA, public resources on health informatics and digital health operations.
7. Industry literature on human-in-the-loop AI in regulated enterprise workflows.
8. Public healthcare operations literature on claims processing modernization and administrative efficiency.