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

# Governance-by-Design Reference Architecture for Agentic Healthcare CRM Using Salesforce Agentforce

**Susil Sahu**

Independent Enterprise Salesforce Architect – Healthcare & Insurance Domain

**Abstract**
Healthcare organizations are increasingly exploring agentic AI to improve service responsiveness, guide users, and support workflow execution. In regulated healthcare environments, however, AI agents cannot be introduced in the same way they might be used in lower-risk domains. This paper presents a governance-by-design reference architecture for agentic healthcare CRM using Salesforce Agentforce. The proposed model introduces autonomy tiers, protected health information minimization, policy-aware orchestration, human oversight boundaries, and audit-ready evidence capture as core design elements. It is intended as a practical architecture pattern for healthcare organizations that want to operationalize agentic AI in a way that is scalable, credible, and defensible.

**Keywords**
Salesforce Agentforce; Healthcare CRM; Agentic AI; Autonomy Tiers; PHI Minimization; Enterprise Governance; Auditability

## 1. Introduction

Healthcare service organizations are increasingly under pressure to improve responsiveness, reduce friction, and deliver better coordination across member service and operational workflows. The arrival of agentic AI introduces new possibilities for contextual guidance, recommendation, and bounded workflow support.

Within regulated healthcare environments, however, AI agents must be governed with greater discipline than in less sensitive domains. Privacy, workflow accountability, and evidence requirements make governance a design concern rather than a downstream compliance concern.

## 2. Problem Context

Healthcare CRM environments involve sensitive member data, policy-driven processes, and actions that may carry service, operational, or financial consequences. An under-governed AI agent can expose unnecessary data, generate misleading recommendations, or trigger actions without the right level of oversight.

This means agentic healthcare CRM must be treated as a governed enterprise architecture problem rather than a generic AI productivity initiative.

## 3. Governance-by-Design Architecture

The proposed model includes six interdependent layers: interaction, context assembly, policy and trust, agent reasoning, orchestration, and evidence. Together, these layers create a practical control framework for useful but bounded AI assistance in healthcare service workflows.

The architecture is designed to embed governance into the operating model from the outset rather than attempting to retrofit governance later.

## 4. Autonomy Tiers Table

- Tier 0: Retrieval and summarization only; no recommendation or execution authority.
- Tier 1: Recommendation support with mandatory human review.
- Tier 2: Constrained execution under narrow policy boundaries and strong logging.
- Tier 3: Broad autonomous execution, generally restricted or avoided in regulated healthcare CRM contexts.

## 5. PHI Minimization and Evidence

Protected health information minimization is a central control principle in this architecture. Agents should retrieve only the information necessary for a given task instead of broad sets of data that increase privacy and compliance risk.

The evidence model records what information was used, what recommendation was produced, what rules applied, who approved the action, and how the workflow concluded. This converts AI activity into an auditable and reviewable enterprise pattern.

## 6. Strategic Value

A governance-by-design architecture helps organizations adopt agentic AI without accumulating governance debt. It improves trust, clarifies allowable behavior, and gives leadership teams a stronger basis for approving new use cases.

For enterprise healthcare platforms, this creates a more durable operating model than experimentation-first deployment.

## 7. Conclusion

Healthcare organizations need a disciplined way to introduce agentic AI into customer and member-facing operations. Governance-by-design provides that discipline by embedding autonomy limits, privacy controls, evidence capture, and human oversight directly into the architecture. This makes responsible innovation more achievable in regulated healthcare environments.

**References**

1. Salesforce, "Agentforce Overview," Salesforce Documentation.
2. Salesforce, "Salesforce Platform Security and Identity," Salesforce Documentation.
3. Salesforce, "Health Cloud and Service Workflows," Salesforce Documentation.
4. Salesforce, "Data Cloud Security and Governance," Salesforce Documentation.
5. JMIR AI and JMIR Medical Informatics, public resources on AI in healthcare operations.
6. AMIA, public resources on health informatics and trustworthy health AI.
7. Industry literature on AI governance, explainability, and model risk management.
8. Public enterprise architecture guidance on auditability, access control, and responsible AI adoption.