SALESFORCE AGENTFORCE • ENTERPRISE AI ARCHITECTURE

The Agentforce Action Blueprint: 5 Custom Actions Your Enterprise Needs Before Turning On AI

Agentforce without guardrails is a liability. Download the Genetrix Agentforce Action Blueprint — 5 production-ready custom actions with system prompts, data constraints, and Apex Invocable structures that make enterprise AI execute safely.

by Genetrix Technology

Published June 10, 2026 • genetrix.tech/blogs

Genetrix Consulting LLC
304 S. Jones Blvd #5631, Las Vegas, NV 89107, USA | genetrix.tech

The Agentforce Action Blueprint: 5 Custom Actions Your Enterprise Needs Before Turning On AI

By Genetrix Marketing • Published June 10, 2026 • https://genetrix.tech/blogs/the-agentforce-action-blueprint-5-custom-actions-your-enterprise-needs-before-turning-on-ai/

Agentforce without guardrails is a liability. Download the Genetrix Agentforce Action Blueprint — 5 production-ready custom actions with system prompts, data constraints, and Apex Invocable structures that make enterprise AI execute safely.

Everyone wants to buy Agentforce. Almost no one has the architecture to support it.

Enterprises are rushing to turn on Salesforce's autonomous AI platform, expecting it to magically resolve support cases, qualify leads, and manage campaigns straight out of the box. The Salesforce demos are compelling. The marketing is compelling. What is not compelling is what happens when you connect an AI agent to a messy CRM and let it run without guardrails.

At Genetrix, we do not deploy AI on hope. We deploy it on strict, documented, tested architecture.

The Harsh Reality of Enterprise AI Deployment

An autonomous agent is only as smart as the strict parameters you build for it. If you plug Agentforce into a CRM with duplicate contacts, unmapped account hierarchies, and inconsistent field values, you are not deploying AI. You are deploying a very expensive chatbot that will confidently hallucinate bad data to your customers — and potentially your prospects.

> **What "Agentforce hallucination" looks like in practice:** An AI agent resolves a support case by recommending a product that was discontinued six months ago. A lead qualification agent tells a prospect their account manager is a former employee. An order lookup agent returns the wrong account's order history because your Contact-to-Account relationships are not clean. All of these are real failure modes we have been brought in to fix.

To make Agentforce execute tasks safely, your architecture team must build custom Actions — binding the AI to specific Apex Invocables, strict Flows, and uncompromising system prompts. You cannot let the AI guess. Every decision point needs a documented boundary.

The Agentforce Action Blueprint: 5 Custom Actions That Actually Work

The blueprint details five battle-tested Agentforce Custom Actions we have built and deployed for enterprise clients. Each one includes the exact system prompt, the data constraints, and the Apex Invocable structure required to make the AI execute safely — without hallucinating, without over-reaching its permissions, and without exposing data it should not touch.

5 Custom Actions Documented in the Blueprint

- **Secure Order Lookup** — Constrains the agent to only surface orders matching the authenticated session's Account ID, with hard stops on cross-account data access
- **Case Escalation Suppression** — Prevents the agent from escalating cases that match a defined suppression list, with audit logging on every suppressed decision

- **Lead Qualification Handoff** — Scores inbound leads against a predefined ICP rubric and routes them to the correct queue without accessing deal history it should not see
- **Campaign Enrolment Guard** — Validates that a contact meets all consent, suppression, and frequency rules before allowing the agent to trigger a campaign enrolment
- **Knowledge Base Scoped Response** — Limits the agent's answer generation to a specific Knowledge Article record type, preventing it from drawing on CRM data outside its approved scope

> **Download the Agentforce Action Blueprint »**
>
> Free PDF · 5 production-ready custom actions · System prompts, data constraints & Apex structures

The Architecture Principle Behind Every Action

Every custom Action in the blueprint is built on the same architectural principle: the AI should be the decision surface, not the data surface. Agentforce's job is to understand intent and select the right action. Every data operation — every lookup, every write, every API call — should happen inside a tightly scoped Apex Invocable or Flow that the AI cannot modify, only trigger.

This separation is what makes enterprise AI safe. The AI decides what to do. A human-written, human-reviewed piece of Apex or Flow determines how it is done and what guardrails apply. The system prompt defines what the AI is not allowed to attempt. All three layers work together.

`Secure_Order_Lookup_SystemPrompt.txt`

```
/* Example System Prompt - Secure Order Lookup Action */
You are an order support agent for Acme Corp.
You may ONLY retrieve orders for the account
associated with the authenticated session.
You must NEVER cross-reference orders across accounts.
If no orders are found, say 'No orders on file' and
offer to connect the customer with a human agent.
Do not speculate. Do not access Case records.
```

Frequently Asked Questions

Do I need Data Cloud for Agentforce to work?

Not for all use cases. Basic Agentforce deployments for case deflection and knowledge base responses can run on standard CRM data. However, for use cases that require a unified customer profile — like personalised campaign enrolment or cross-channel lead qualification — Data Cloud is strongly recommended as the data layer. Without it, the agent is limited to whatever structured data exists in your CRM objects.

What is the difference between an Agentforce Action and an Agentforce Topic?

A Topic defines what the agent is for — its domain, scope, and the types of conversations it should handle. An Action is what the agent can do within that scope — the specific operations it can trigger, the data it can access, and the systems it can interact with. The blueprint focuses on Actions, because that is where most enterprise implementations fail: the Topics are set up correctly, but the Actions are too broad or too weakly scoped.

How do we prevent Agentforce from accessing data it should not?

Three mechanisms working together: the system prompt (instructional boundary), the Apex Invocable's query scope (technical boundary), and the Agentforce Permission Set assigned to the agent's running user (platform boundary). All three must be configured correctly. Relying on the system prompt alone is insufficient for enterprise-grade data security.

How long does it take to build and deploy one custom Agentforce Action?

For a well-scoped Action with a clear data model, an experienced developer can build and test the Apex Invocable in 1–2 days. System prompt engineering and UAT typically add another 1–3 days depending on the

complexity of the edge cases you need to handle. The five Actions in the blueprint represent roughly 2–3 weeks of total build and testing time across a small team.

Ready to Deploy Agentforce the Right Way?

Genetrix designs and implements enterprise Agentforce deployments with the strict guardrails, tested Actions, and documented system prompts that enterprise AI requires. If you are planning an Agentforce rollout or need an architecture review of an existing deployment, our team is ready.

Get in Touch with Genetrix »