Agentforce for Salesforce Teams: Use Cases, Governance, and Guardrails

CHEATSHEET

How to structure permissions, define clear topics, and enforce safety boundaries in your Salesforce org.

Deploying Agentforce without setting strict Topics and Guardrails is the fastest way to get your Salesforce AI project pulled by compliance. Here is your governance playbook.

1 The 3 Pillars of Agent Behavior

Every Agentforce agent is governed by three critical layers. If you don't structure these properly, your agent can access records it shouldn't.

- Agent Topics: The domains of work (e.g., Case Management) — the agent's job description.
- Agent Actions: The specific tools/operations (e.g., querying records, sending emails).
- Guardrails: Explicit, unbreakable rules that stop agents from acting outside their scope.

2 High ROI, Low-Risk AI Use Cases

Start your deployment with high-volume, repetitive tasks where the cost of an error is low. Avoid complex financial transactions in your initial launch phase.

- Case Deflection: Deflecting routing queries via self-service portal tools.
- Knowledge Retrieval: Quick internal Q&A for sales and onboarding support.
- Lead Qualification: Filtering incoming leads before assigning to SDRs.
- Order Lookups: Checking delivery status in Commerce Cloud without manual search.

3 Hardening Your Agent Governance

Prevent your autonomous AI agent from exposing draft data or bypassing internal checks. Always configure these guardrails before going live.

- Strict FLS: Ensure agents respect the running user's Field-Level Security.
- Object Restrictions: Isolate Case agents from accessing sensitive Opportunity records.
- Manual Handoffs: Always define a clear pathway to route complex queries to live agents.
- Approval Controls: Prevent agents from bypassing multi-step approval workflows.