Prompt Builder

Salesforce Certified AI Specialist

KNOWLEDGE CHEATSHEET

REUSABLE GENERATION

Choose Template → **Ground** Context → **Preview** Validate → **Activate** Use

01 Prompt Builder and reusable templates

Choose the template for the outcome

- **Prompt Builder** creates reusable instructions for generating content in Salesforce workflows. Choose the template type from the destination and required inputs.
- **Sales Email** templates produce personalized outreach using sender, recipient and relevant CRM context. Ground product recommendations and follow-ups in customer interests and actual product information.
- **Field Generation** produces content for a supported record field, such as an account or guest summary. Configure the target object, field and record-page experience.
- **Flex** supports custom generation tasks with multiple configured inputs, including unrelated records. Use it for a newsletter, recommendation or summary that does not fit a specialized template.
- **Record Summary** templates control record summarization where supported. Match the object and summary experience to the business need; a summary template and a field-population template serve different destinations.

Write a clear generation contract

- Specify the **task, role, audience, tone and output format**. State which supplied facts may be used, what the response must include, and how to handle missing information.
- Use **grounding** for factual customer context and examples for desired style. Examples of tone do not establish prices, contract terms or other business facts.
- Separate **instructions from retrieved data** with clear labels and delimiters. Quotes or braces organize a prompt; punctuation alone does not enforce compliance or prevent injection.
- Request a bounded response: relevant facts, concise length and a clear next step. Tell the model to acknowledge unavailable information instead of inventing values.
- **Iterate with representative records** and user feedback. Evaluate factual accuracy, completeness, tone, privacy and usefulness; a fluent response can still be wrong.

Prompt Builder / Grounding & Models

Salesforce Certified AI Specialist

KNOWLEDGE CHEATSHEET

REUSABLE GENERATION

Choose Template → **Ground** Context → **Preview** Validate → **Activate** Use

Preview, resolve and activate

- Configure inputs and grounding resources, select the model, then **save and preview** with sample records. Missing required preview inputs can prevent generation.
- **Resolution** is the prompt after dynamic resources have been resolved. **Response** is the model-generated output. Inspect both to separate grounding problems from generation problems.
- For a poor summary, check **missing, stale, incorrect or inaccessible source data** before rewriting instructions. Compare resolved values with the selected record and intended related records.
- Create an editable **new template** when adapting standard content; use a **new version** to revise an existing custom template. Test the revision and activate the intended version for runtime use.
- **Token usage varies by record** because related lists, long text and generated output vary. Limit retrieved records and fields, summarize large inputs, and allow room for the response in the model context.

Expose prompts in the flow of work

- **Dynamic Forms** supports adding field-generation templates to eligible Lightning record pages. Verify object and field support, associate the template with the field, and activate the page configuration.
- For mobile field generation, check the supported **Dynamic Forms on Mobile** configuration and the user experience on the actual device.
- **Prompt Template Manager** supports creating and managing templates; **Prompt Template User** supports executing them. Feature access still depends on licensing, org enablement and data permissions.
- A **screen flow** can collect an order number, retrieve the order and details, invoke an active template through a Flow action, and display its response. A background workflow can invoke a template without a screen.
- For application integration, use supported **REST or Apex prompt-template invocation**. Authenticate the caller, supply the required inputs, handle generation failures, and validate output before saving it.

02 Grounding, Data Cloud and model choices

Supply relevant CRM context

- **Record merge fields** insert values from the selected record at runtime. Personalization depends on accurate data and the executing context having access to the needed fields.
- **Related-list grounding** supplies child records, such as an account's opportunities. Check the supported relationships and related-list layout fields rather than assuming every relationship or activity list is available.
- Use a **flow or Apex** for selective retrieval, sorting, aggregation or transformation. A requirement for the three newest opportunities needs an explicit selection rule, not an arbitrary list of opportunities.

Grounding & Models

Salesforce Certified AI Specialist

KNOWLEDGE CHEATSHEET

CHOOSE THE RIGHT LAYER

- **Grounding** — Runtime facts
- **Prediction** — Scores
- **Generation** — Content
- **Tuning** — Model behavior

- Retrieve **only necessary fields and records**. Define behavior for empty lists, missing fields and ambiguous matches; excessive context can increase token cost and bury relevant facts.
- **Grounding is not fine-tuning**: grounding adds runtime evidence to a prompt; fine-tuning changes model behavior through training. Neither removes the need to validate the generated response.

Extend grounding with Flow and Apex

- A **template-triggered prompt flow** gathers and formats data for a template. Its **Add Prompt Instructions** element appends text and resource values to the prompt material returned by the flow.
- Use flow logic to combine CRM records with **external service responses**, such as nearby events or fulfillment status. Validate the returned data and format it before inserting it into the prompt.
- **Invocable Apex** can prepare complex grounding data, including serialized JSON. Use proper JSON serialization; instructions to return JSON do not by themselves guarantee valid structured output.
- **Data Cloud** can provide unified customer context and activity data. Supported enrichment related lists or a grounding flow can bring relevant unified-profile information into a CRM-based prompt.
- Distinguish **a flow supplying prompt context** from **a flow invoking a prompt**: the former returns grounding instructions; the latter consumes generated output as a workflow result.

Separate prediction, generation and tuning

- **Predictive AI** estimates outcomes or scores, such as churn risk, lifetime value or product propensity. **Generative AI** drafts text, summaries and personalized messages.
- Use a **prediction or explicit business rule** to choose a segment or next action when the choice must be explainable. A generator can then write the message; its narrative is not evidence of why a predictive model made a decision.
- **Model Builder / Einstein Studio** supports configuring supported models and connecting external models. A fine-tuned external LLM needs a supported connection, credentials and an accessible inference endpoint.
- For an externally trained predictive model, such as **XGBoost on Databricks**, expose a supported serving endpoint and configure the connection and input/output mapping. Arrange ingestion of inference results into the required **data model object (DMO)**.
- **Temperature** adjusts sampling variability; **frequency penalty** discourages repeated tokens; **presence penalty** discourages reuse of already-present tokens. Available controls depend on the model and configuration.
- Use **Model Playground** to compare supported model configurations. Copying prompt text does not ensure identical model settings or results; record the chosen model, parameters and output constraints with the implementation.

Agents & Actions

Salesforce Certified AI Specialist

KNOWLEDGE CHEATSHEET

CONVERSATIONAL WORK

Request Understand → **Plan** Select actions → **Execute** Flow or Apex → **Respond** Use results

03 Conversational agents and actions

From a request to authorized work

- **Agentforce** provides a conversational interface to retrieve information, generate content and perform configured business tasks. Available actions and accessible data define what an agent can accomplish.
- The **planner or reasoning engine** interprets the user's request and conversation context, selects suitable actions and determines their sequence. A model plans; the referenced Flow, Apex or other supported action performs the operation.
- **Standard actions** cover supported tasks such as record summaries, email drafting, finding similar opportunities and answering questions with Knowledge. Prefer a suitable standard action before building custom logic.
- When a request is unclear or data cannot be found, ask for **clarification** or explain the limitation. Do not fabricate a record, assume an identifier, or imply that an unsuccessful action completed.

Build actions the agent can distinguish

- A **custom action** can reference a supported autolaunched flow, invocable Apex or prompt template. Reuse an existing ERP lookup flow to retrieve live fulfillment status through conversation.
- Define a specific **action name and instructions**: what the action does, when it should run, and its boundaries. Distinguish similar summary actions by purpose and required context.
- Describe **inputs** precisely: meaning, data type, format and where the value comes from. Require necessary inputs and collect missing values before execution.
- Describe **outputs** and how the agent should use or display them. Return useful data and explicit success or failure information; an output being returned does not mean every field should appear in the conversation.
- State **dependencies and sequencing**, such as looking up an account before querying its orders. Assign the action to the agent's applicable topic and configure confirmation where required for record changes.

Agents & Actions / Trust & Rollout

Salesforce Certified AI Specialist

KNOWLEDGE CHEATSHEET

CONVERSATIONAL WORK

Request Understand → **Plan** Select actions → **Execute** Flow or Apex → **Respond** Use results

Test selection and execution separately

- Use **Agent Builder** to test varied real-world utterances before activation. Include paraphrases, competing intents, missing values and requests outside the action's scope.
- For the **wrong action**, inspect topic scope, action assignment, names and instructions. Remove overlapping generic wording and make each action's intended request distinct.
- For the **right action with wrong results**, inspect selected action, input values, returned outputs and the underlying flow or Apex. Test the reference implementation independently.
- A request to **delete or update records** needs an available, configured action plus appropriate execution permissions. Natural-language understanding does not create unsupported capabilities or bypass access checks.
- Use **conversation event logs, utterance analysis and agent analytics** to find repeated misunderstandings, action failures and adoption patterns. Enable relevant logging deliberately and protect captured conversation data.

04 Trust, security and controlled rollout

Protect the request and response

- **Einstein Trust Layer** combines secure retrieval, grounding, privacy controls and auditability around model use. Apply the controls supported by the particular feature and execution path.
- **Secure data retrieval** respects applicable access controls when obtaining CRM context. Prompt instructions cannot grant record access or override field-level security.
- **Data masking**, where supported, replaces selected sensitive values before external model processing; **demasking** restores mapped values on the response journey. Configure sensitive entities or fields according to policy.
- **Agent actions have different masking behavior** from direct template use: do not assume a prompt invoked through an agent is masked because the same template is masked when used directly.
- **Zero data retention** concerns third-party model-provider retention and training use. It does not mean Salesforce stores no records or audit data, or that sensitive input can be sent without a privacy assessment.

Trust & Rollout / Sales & Service

Salesforce Certified AI Specialist

KNOWLEDGE CHEATSHEET

CONTROLLED AI USE

- **Access** — Retrieve safely
- **Privacy** — Apply controls
- **Safety** — Check output
- **Audit** — Monitor use

Distinguish safety signals from correctness

- **Prompt defense** helps resist jailbreaking and prompt injection. Treat retrieved documents, user text and external responses as data, not as authority to change system instructions.
- **Toxicity scoring** assesses harmful-language risk. Higher category scores indicate greater risk; a safety score does not establish factual accuracy, fairness or legal compliance.
- **Hallucination** is unsupported generated content. Relevant grounding, clear constraints and human review reduce risk, but neither masking nor a low toxicity score guarantees a correct answer.
- After changing **masking configuration**, test both protection and response usefulness with representative sensitive values. Verify what actually reaches the model rather than judging only the visible final response.

Audit and deploy with least privilege

- Enable **generative AI audit and feedback collection** for the supported org configuration. Audit data in Data Cloud can include resolved or masked prompts, responses and trust signals such as toxicity scores.
- Use the **AI audit trail** to verify masking and investigate generation behavior. Filter the relevant trust signals in reports; ordinary platform debug logs and Setup Audit Trail serve different purposes.
- Separate **creator, executor and data-administrator permissions**. Grant only the required feature access, record access and field access; a shared profile does not imply identical permission-set assignments.
- Configure and test in a **sandbox**. Deploy supported metadata and dependencies, including fields, flows, prompt templates, quick actions and record pages; separately verify org-level enablement, credentials and user assignments.
- Before rollout, test **representative user contexts**, negative cases and integration failures. After rollout, monitor feedback, usage and quality, then revise the relevant template, action or grounding logic.

05 Sales and service applications

Draft outreach and guide deals

- **Einstein Sales Emails** and email-drafting actions create personalized drafts from relevant customer context. Review names, claims and commercial commitments before sending.
- **Draft or Revise Sales Email** supports creating and refining outreach. A tone adjustment such as **Make Less Formal** changes style; it does not supply missing customer facts.

Sales & Service

Salesforce Certified AI Specialist

KNOWLEDGE CHEATSHEET

MATCH THE BUSINESS TASK

- **Email** — Draft outreach
- **Calls** — Extract insights
- **Replies** — Ground in Knowledge
- **Wrap-up** — Summarize work

- If a draft contains **unresolved placeholders**, check record completeness, input mapping and field permissions. More elaborate instructions cannot recover data the executing context cannot read.
- **Find Similar Opportunities** supports learning from eligible comparable deals. **Create Close Plan** provides guidance for progressing an opportunity; confirm suggestions against the actual deal and business process.

Use the right call capability

- **Einstein Conversation Insights** analyzes supported recorded calls and transcripts for relevant mentions and coaching signals. Connect the recording source and assign the required access.
- **Call Insights** highlight configured products, competitors and phrases. Configure the business vocabulary and verify transcription quality before relying on mention counts.
- **Call Summaries** condense a conversation into key points and next steps. Use them to recover missed details from long voice or video meetings; verify significant deal terms against the recording or transcript.
- **Call Explorer** supports conversational questions about call content. Distinguish asking a question about a meeting from reviewing an automatically generated summary or tracked keyword insights.

Ground replies and reduce wrap-up work

- **Einstein Service Replies** drafts responses during supported service conversations. Email reply generation addresses service email drafting; choose the capability for the actual channel.
- **Service AI Grounding with Knowledge** supplies relevant article content for responses. Publish and maintain accessible product guides and tutorials as Knowledge articles, and review supporting references in the generated reply.
- **Knowledge Recommendations** suggests relevant articles; **Reply Recommendations** suggests reusable replies. Distinguish recommending existing content from generating a new grounded response.
- **Record or case summaries** help an agent understand an existing case. **Work Summaries** reduces interaction wrap-up effort by generating information such as issue, resolution and summary from conversation context.
- For **Enhanced Messaging Work Summaries**, configure storage fields, a quick action that updates them, and the **Wrap Up** component on the Messaging Session record page. Include these dependencies in the deployment plan.
- Agents should **review suggested replies, summaries and field values** before committing them. An omitted fact may reflect missing conversation context or inaccessible grounding; a summary is not a substitute for the original record.