

Dynamic Contexts: A Data-Driven Context Framework Applied to CRM

March 2026

Introduction

AI agents deployed in customer relationship management are only as effective as the framing information they receive, and when they receive it. Customer data, however rich, is not enough: it tells you who the customer is, not how to speak to them, what market context they are operating in, or what the company is authorized to offer them. This is precisely the role of context: to complement data so that the agent can produce responses that are tailored, coherent, and effective.

Context is not a static block of information written once and never revisited. Some of its layers are stable and rarely change; others update in real time as market dynamics or available inventory shift; others still build progressively from past interactions. Above all, context maintains a strong structural relationship with customer data: data does not constitute context, but it often activates it. A customer's address triggers the retrieval of the right local market context. Its financial profile filters the applicable offer rules. Its past behavior guides the search through the learning memory. Making this connection between data and context explicit, understanding how one feeds the other without becoming the other, is one of the central challenges this document takes on, through a generic framework and concrete examples.

This document aims to define and structure the notion of data-driven context within agentic AI in general, then to apply it specifically to the CRM domain. It begins with a generic framework organized around nine context layers, grouped into three families according to their degree of stability: stable contexts that establish the advertiser's identity and constraints; dynamic contexts that anchor the agent in the reality of the market and available offering at the moment of interaction; and inferred or learned contexts that allow the agent to become cumulatively more relevant. It then details how each layer is practically injected into the prompt architecture, and how customer data interacts with these contexts, not replacing them, but activating, filtering, or orienting them. The framework is illustrated through a concrete case in the real estate development sector, across four contrasting scenarios (chatbot, investor email, premium primary residence email, first-time buyer email), showing how the same development, addressed to different profiles, produces radically distinct communications once context is properly architected.

This document does not examine the production deployment of such generative AI use cases, nor the technical solution to build an application capable of producing such a context. Similarly, it does not address latency, cost, or security considerations.

Table of content

Introduction.....	2
Table of content	3
Defining Context Types.....	4
Stable Contexts.....	4
Dynamic Contexts.....	5
Inferred or Learned Contexts.....	5
Prompt Architecture.....	6
Stable Contexts → System Prompt	6
Dynamic Contexts → Injection at Execution Time (RAG & Tools).....	6
Inferred or Learned Contexts → In-Interaction Construction + Memory.....	7
The Full Context Picture.....	7
The Relationship Between Customer Data and Context	9
Customer Data as a Context Personalization Filter.....	9
Customer Data as a Search Vector in Memory.....	9
What Dynamic Contexts Change About the Architecture	9
Implications for Data Governance	9
A Two-Tier Architecture for Producing This Type of Prompt	11
The Orchestrator: Deterministic Workflow or LLM?	11
What Most Production Systems Do.....	12
Reference Architecture for Dynamic Contexts.....	13
The Real Estate Development Examples.....	15
What the CDP Can Provide in This Industry	16
Example 1: Chatbot.....	16
Example 2: Relationship Marketing Email for an Investor.....	21
Example 3: Relationship Marketing Email for a Family Purchase.....	28
Example 4: Relationship Marketing Email for First-Time Buyers.....	35
Comparison Across the Three Email Examples.....	43
Conclusion	44
Appendices	45
Sample Prompts for Example 1.....	45
Sample Prompts for Example 2	49

Defining Context Types

Why distinguish between types of contexts rather than feeding the agent a single list of information? Because not all framing information behaves the same way. Different pieces of context have different lifespans, come from different sources, are injected into prompts through different mechanisms, and fall under different maintenance responsibilities. Conflating what is permanent with what is real-time, or what is prescribed with what is learned, leads to fragile architectures: either the system prompt becomes overloaded with information that should have been injected dynamically, or elements that should remain stable get lost mid-interaction.

The nine-layer distinction follows a logic of functional granularity: each layer corresponds to a distinct dimension of what the agent needs to know in order to respond appropriately. Who is the advertiser and how should it speak? What rules govern it? What limits has it set for itself? In what market is it operating right now? What can it concretely offer? Through which channel and at what moment is it speaking? What is its objective in this specific interaction? Who is the customer on the other side, and what is their likely intent? What has been learned from similar past interactions? Each of these nine questions calls for a different type of answer, from a different source, with a different update cadence. Grouping them into an undifferentiated list would mean treating a brand principle the same way as a product inventory or a learned behavioral pattern, three things that share nothing: not their source, not their update cadence, not their ownership.

These nine layers are organized into three families according to their degree of stability:

- **Stable contexts** form the permanent foundation of the agent: they describe what the advertiser is, what it is bound by, and the limits it sets for itself. They change rarely, they belong to brand, legal, and compliance teams, and they live in the system prompt.
- **Dynamic contexts** anchor the agent in the reality of the moment: market conditions, available inventory, active channel. They are retrieved at execution time via tools and connectors and depend directly on the quality of the advertiser's information systems.
- **Inferred or learned** contexts are the most complex: they combine initial injection with progressive construction, depend on customer data and accumulated memory. These are the ones that make the difference between a capable agent and a truly effective one.

We anticipate that the 3 families will apply to all industries (B2B and B2C) and functions, while the context layers inside each of them may vary.

The nine context types are described below.

Stable Contexts

Live in the system prompt, change rarely, maintained at the global level.

1. Brand & Positioning Context

The voice and tone, values, language to use or avoid, the core promise, positioning versus competitors. This is the stable personality embodied by the agent at all times.

2. Regulatory Context

The legal rules that govern the advertiser in its sector and geography: what the agent can and cannot state, mandatory disclosures, topics to avoid entirely. Financial services, healthcare, insurance, and alcohol all operate under different constraints. An error here creates direct legal exposure for the advertiser.

3. Internal Compliance Context

The framework the advertiser sets for itself: internal communication policies, codes of conduct, CSR commitments, sensitive topics specific to the brand or its history. What the company has decided not to do, even where the law does not prohibit it.

Dynamic Contexts

Retrieved at execution time via tools or updated knowledge bases.

4. Market Context (General & Local)

The macro dynamics of the sector, seasonality, competitive pressure, and local specifics: regional behaviors, competitor offers active in a given territory, local events that shift intent or timing.

5. Business & Offer Context

The current state of the catalog, inventory levels, current commercial priorities, terms of sale, eligibility constraints. Must be maintained in near-real time to avoid erroneous proposals.

6. Channel & Moment Context

The conventions of the active touchpoint (email, chat, SMS, app), format constraints, time of day or year, level of interactivity available.

Inferred or Learned Contexts

Built progressively. Make the difference between a capable agent and a truly effective one.

7. Agent Objective & Mission Context

What the agent is mandated to do in this specific interaction: convert, retain, collect information, handle dissatisfaction, or conversely, explicitly not sell. It can be recalibrated during the interaction based on incoming signals.

8. Relational & Intentional Context

The interpretation of customer data: where is the customer in their journey, what is their probable state toward the brand, what intent lies behind the current interaction. This is what transforms cold data into actionable signal.

9. Learning & Performance Context

What has worked or failed in similar situations: formulations that convert, frequent objections and their responses, early churn signals. This is the context that enriches over time and makes the agent cumulatively more relevant.

The further down the list, the more **alive, situated, distinctive and interaction-specific** the context but also the harder they are to build and maintain. Stable layers protect and frame. Dynamic layers anchor in reality. Inferred layers are where differentiated value is created.

Prompt Architecture

Stable Contexts → System Prompt

These layers are written once, versioned, and systematically injected at the beginning of every prompt as the agent's base instruction. They do not change from one interaction to the next.

Brand & positioning is written as persona instructions: "You are the assistant for [Brand]. You communicate in a warm but expert tone. You never say X; you always prefer Y. Your core promise is..." This is the permanent identity foundation.

Regulatory is injected as explicit, non-negotiable rules: "You can never claim this product cures anything. You must always mention... in the case of a commercial proposal. You refuse any request that would lead you to..." The agent must not interpret these rules, it must apply them literally.

Internal compliance follows the same logic as regulatory, with a nuance: these rules may be slightly more contextual. *"If this topic is raised, redirect to..."* or *"You never make a direct comparison with [competitor]."* Often written as a list of guardrails.

→ **Key architectural point:** these three layers together form the master system prompt. It is maintained by brand, legal, and compliance teams, not by the teams who operate the agents day to day. It is alive, versioned and governed.

Dynamic Contexts → Injection at Execution Time (RAG & Tools)

These layers cannot be written in advance. They are retrieved dynamically just before or during response generation.

General & local market is retrieved via tools connected to competitive intelligence feeds, geolocated databases, and market signal aggregators. The agent receives a contextualized snippet along the lines of: "Current market context: strong promotional pressure from [competitor] in the Southwest region over the past 15 days. Seasonality: back-to-school period, active demand." The agent does not interpret the market: a specialized tool pre-digests and indexes the information for it. It can be generated through an agentic process.

Business & offer is injected via a real-time call to the product or catalog base, ideally through a dedicated tool the agent can query. The agent does not know the catalog by heart; it consults it just when it needs it, like a sales representative checking inventory before promising a delivery date. This prevents hallucinations about offers and pricing.

Channel & moment is partly in the system prompt (the fixed rules of the channel) and partly injected dynamically at this level (the time, the day, the commercial period). An application wrapper passes these variables to the agent: *"Channel: SMS. Max length: 160 characters. Moment: Sunday evening. Period: 3 days before end of promotion."*

→ **Key architectural point:** these three layers depend on the quality of the tools and connectors made available to the agent. An agent without the right tools will fill those gaps by inventing or overgeneralizing, and this is precisely where most relevance errors originate.

Inferred or Learned Contexts → In-Interaction Construction + Memory

These layers are the most complex to architect because they combine initial injection with progressive construction.

Agent objective & mission is injected at the start of the interaction via what is called a task prompt, distinct from the system prompt: *"In this interaction, your objective is to retain a customer who has just expressed an intent to cancel. You only propose a commercial offer as a last resort, after understanding the reason."* This objective can be dynamically recalibrated if the agent detects that the situation is evolving, especially for ongoing interactions such as chatbots. This is where genuinely agentic behavior begins.

Relational & intentional is the most sophisticated layer. It combines three sources:

- customer data injected via RAG (history, status, recent behavior).
- a deduction made by the agent itself from the first signals of the interaction (*"the customer opens with a complaint; their tone is tense"*).
- and potentially an external scoring module that pre-qualifies the relational state before the agent takes over.

In practice, a structured context summary is often built to synthesize these signals: *"Customer: high value, 4 years tenure, unresolved incident 8 days ago, last interaction negative. Probable intent: escalation."*

Learning & performance is the long-term memory layer. It is injected via a vector database that stores past interactions, their outcomes, and the patterns that worked. Before responding, the agent (or an orchestrating agent) queries this database to retrieve similar cases: *"In analogous situations, the formulations that converted best were... The most frequent objections are... Early abandonment signals are..."* This is what transforms the agent from a system that resets with every interaction into one that compounds its relevance over time.

→ **Key architectural point:** these three layers require a memory infrastructure (short-term within the conversation context, long-term in an external database) and often an orchestrating agent that prepares and structures context before the surface agent speaks. The work is not done by a single agent: context preparation is separated from response execution.

The Full Context Picture

Schemas such as the one below represent the structured content transmitted to the orchestrator. They are not executable code. They illustrate the nature and organization of the information assembled in the prompt before the agent speaks. Each line corresponds to an active piece of context, with its source indicated on the right.

```

SYSTEM PROMPT (stable)
├─ Brand & positioning
├─ Regulatory
└─ Internal compliance

+ DYNAMIC INJECTION (execution)
├─ Market & local           ← intelligence tools
├─ Business & offer        ← catalog connectors
└─ Channel & moment        ← application wrapper

+ PROGRESSIVE CONSTRUCTION (inference + memory)
├─ Objective & mission     ← task prompt + recalibration
├─ Relational & intentional ← customer RAG + inference
└─ Learning                ← long-term vector database

```

The quality of a marketing agent is not primarily a function of the language model chosen. It is a function of how well these nine layers are fed, maintained, and orchestrated. Some examples are provided in the Real Estate Development scenarios below.

The Relationship Between Customer Data and Context

The line between "context" and "customer data" is not fixed. It is, rather, a **triggering relationship**. Customer data does not replace context, but it **activates and personalizes certain layers**.

Customer Data as a Context Personalization Filter

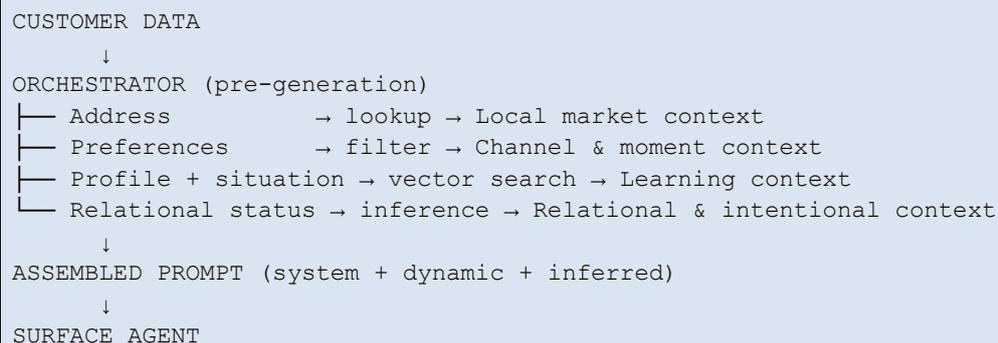
The customer's channel and moment preferences do not define the channel rules. They select which set of rules to apply. The channel & moment context exists for all possible channels. The customer preference tells the agent: "use this subset, not the others." The data filters, it does not construct.

Customer Data as a Search Vector in Memory

For the learning & performance layer, customer data (profile, segment, behavior, current situation) serves as a vector to query the long-term memory base. The system does not retrieve the full accumulated learning: it retrieves cases **similar to this customer in this situation**. The data qualifies similarity and orients the search.

What Dynamic Contexts Change About the Architecture

This dependency introduces a **mandatory intermediate step** between customer data and the final prompt construction: an orchestration layer that, before any response generation, uses customer data to activate, filter, and query the right contexts.



The orchestrator does not generate responses. It has a single responsibility: assemble the most relevant context possible before handing off to the agent that speaks to the customer. This **separation between preparing and executing** is what makes the system robust and maintainable.

Implications for Data Governance

This framework also reveals a frequently underestimated requirement: for the data-dependent context layers to function, customer data must be **clean, structured, and accessible in real time**. A missing address prevents the activation of the market context.

An unrecorded channel preference forces a **generic fallback**. An incomplete history impoverishes the memory search.

The quality of injected context is therefore a **direct indicator of the quality of customer data**. Agentic systems surface every data gap that a human advisor would have papered over with intuition.

A Two-Tier Architecture for Producing This Type of Prompt

To produce this type of prompt, the AI system operates across two tiers:

1. An **Orchestrator** responsible for reasoning about what to inject in order to build the prompt.
2. A **Surface agent** that reasons about what to say based on the orchestrator's instructions and generates individualized content.

Although the Surface agent will always be an LLM, the technology top use for the Orchestrator is debatable. Various options are listed below.

The Orchestrator: Deterministic Workflow or LLM?

This is the central question of agentic architecture. Both solutions are possible depending on design choices. The two approaches involve very different trade-offs.

A deterministic workflow as orchestrator

In its simplest form, the orchestrator is a conventional program (Python, n8n, Make, or a custom service) that executes a hardcoded sequence of steps:

1. Read the incoming message
2. Extract entities (city, intent, segment)
3. Call the market intelligence API with the extracted city
4. Query the catalog database with the profile parameters
5. Retrieve the N most similar cases from the vector database
6. Assemble the prompt from a template
7. Call the LLM

Each step is deterministic and traceable. If "Denver" appears in the message, the "Denver" API is called. If the segment is "investor," the "investor" persona is activated. No ambiguity, no surprises. This is the preferred approach in production when the case space is well-defined and reliability is critical.

- **Advantages:** predictable, auditable, fast, cost-controlled, easy to debug.
- **Limitations:** it can only handle anticipated cases. An ambiguous message, a situation outside the template, and it either fails or produces a generic fallback.

An LLM as orchestrator

In a more sophisticated architecture, the orchestrator is itself a language model, often called an orchestrating agent or planner. The situation and available tools are described to it, and it decides which tools to call, in what order, and how to interpret the results.

A sample prompt of this type appears in the appendix.

The LLM orchestrator reasons, calls tools, interprets results, handles ambiguities, and produces the assembled prompt. This is what frameworks such as LangGraph, CrewAI, and the tool-equipped agents from Anthropic and OpenAI do.

- **Advantages:** flexible, capable of handling unanticipated cases, can reason over ambiguous signals, improves with the underlying model.

- **Limitations:** less predictable, more expensive (two LLM calls instead of one), harder to audit, higher latency.

What Most Production Systems Do

In the real world, a hybrid architecture is most common:

- The skeleton is a deterministic workflow (extraction steps, API calls and data retrieval, prompt structure)
- Certain nodes in the workflow call an LLM for decisions that require judgment (interpreting an ambiguous intent, choosing between two personas, assessing whether a history triggers a specific relational context). In these cases, the LLM is responsible for building the context summary from the workflow results.

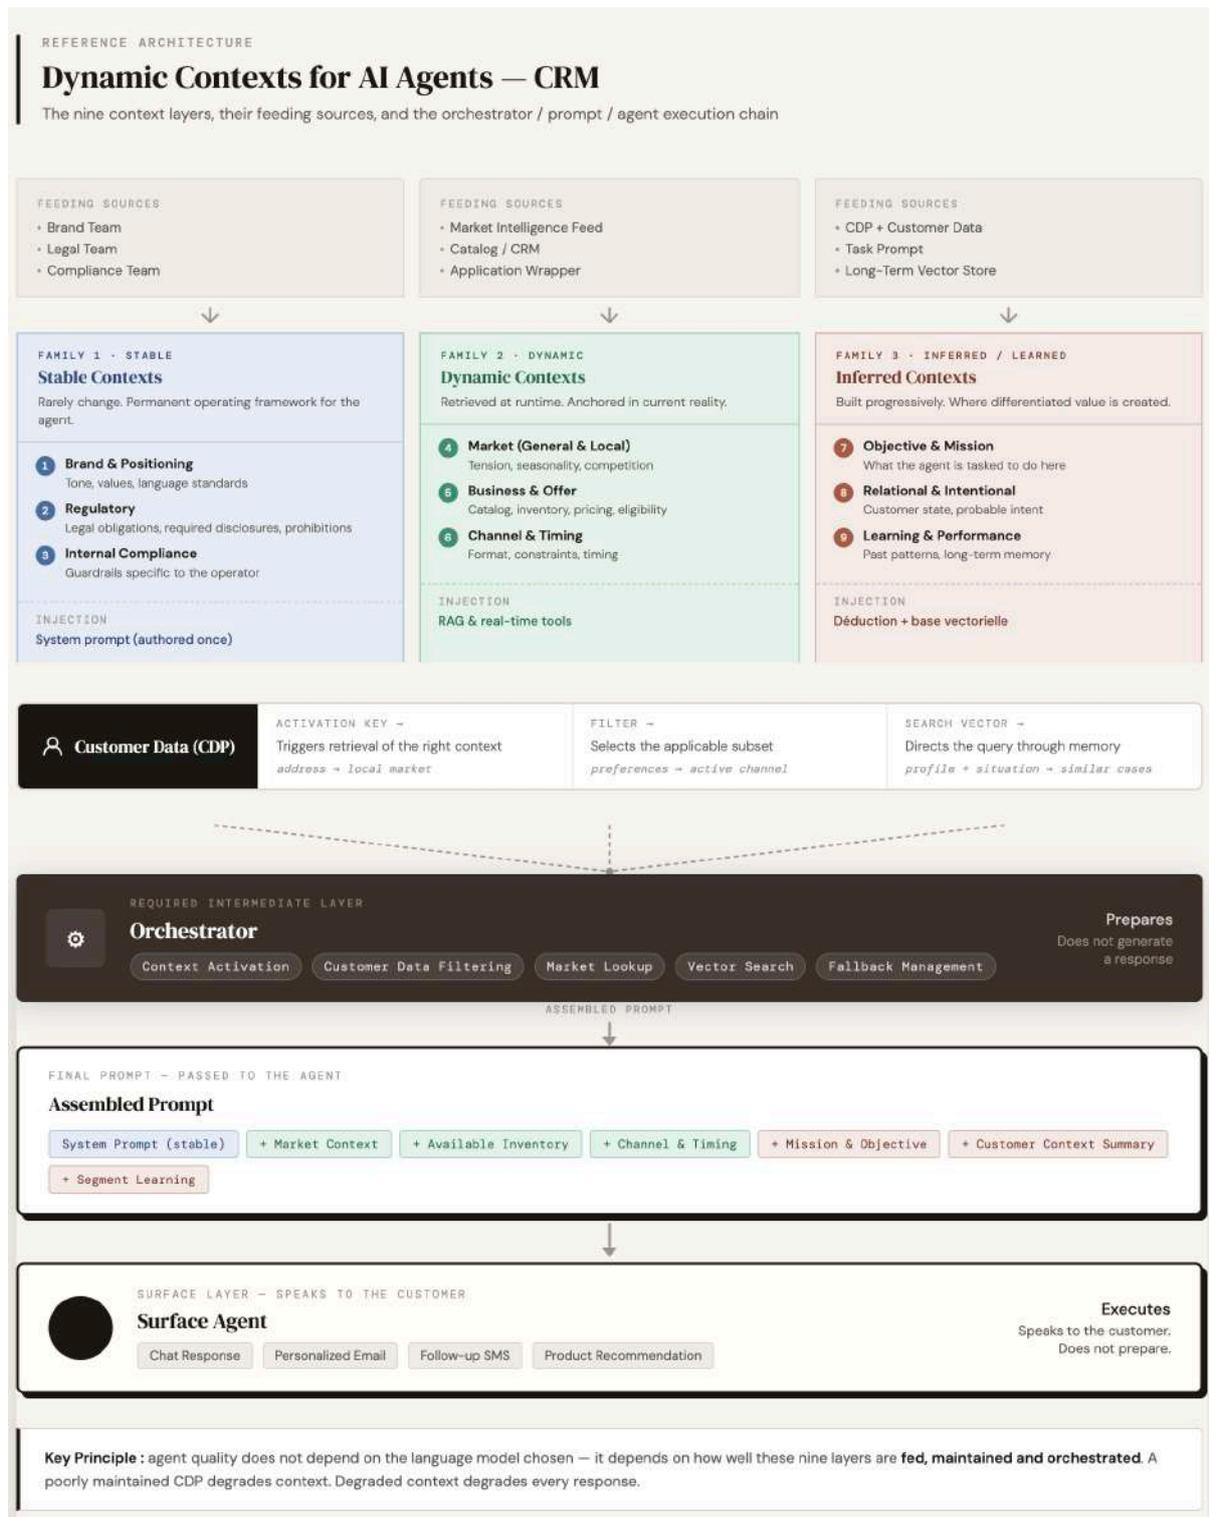
The balance between workflow and LLM is the calibration point that each team must adjust according to four criteria:

1. Risk tolerance (an LLM orchestrator can fail in unexpected ways)
 2. Expected response latency
 3. Infrastructure cost (in CRM use cases with millions of contacts, costs can escalate rapidly)
 4. The real complexity of the cases to be handled

This balance also determines how fallback rules are documented: in a deterministic workflow, they are coded in. In an LLM orchestrator, they are injected as instructions. In both cases they must exist, but their governance is very different.

Reference Architecture for Dynamic Contexts

Dynamic Contexts for AI Agents: CRM



The diagram above presents a global architecture for context construction:

- Feeding sources:** three columns at the top show who produces each family (brand/legal/compliance teams for stable; intelligence feeds + catalog for dynamic; CDP + vector database for inferred).

2. **The nine layers:** the three color-coded families (blue / green / earth tones) with their numbering 1–9, short descriptions, and the injection mechanism noted at the bottom of each column (system prompt / real-time RAG / inference + vector).
3. **Customer data:** a horizontal band crossing all three families, showing its three distinct roles (activation key, filter, search vector) with a concrete example for each.
4. **The execution chain:** the three columns converge on the Orchestrator (which assembles without speaking), then the Assembled Prompt (with its color-tagged layers), then the Surface Agent (which executes). The *prepare / execute* separation is made visually explicit.

The Real Estate Development Examples

The following examples illustrate the application of this methodology to the case of a real estate developer. This sector was chosen because it concentrates nearly all the complexities an agentic CRM system can face. Decision cycles are long (several months), the emotional dimension is high (real estate purchase as a life event), regulation is dense (new construction purchase agreements, FHA, cost segregation, disclosure obligations), local markets are highly heterogeneous (Denver is not Minneapolis), and customer profiles differ radically in their motivations, constraints, and expected register.

Four scenarios are developed:

1. A **chatbot** (unknown visitor, zero history, progressive acquisition of data)
2. An **email** for an **investor** (Michael Chen, warm lead)
3. An **email** for a **premium family** purchase (Alice Morgan, advanced selection phase)
4. An **email** for **first-time buyers** (Tyler & Emma, post-in-person-visit)

Each unfolds through the same steps: CDP reading if possible, activation of each context layer, documented fallback decisions, and assembled prompt. Taken together, they demonstrate how the same development, when addressed to different profiles, yields communications that share virtually no paragraph once context is properly architected.

What the CDP Can Provide in This Industry

Some of the examples draw on data from the CDP. Before running the flow, it is worth being precise about what a real estate CDP would contain:

IDENTITY & LOCATION

- |— Declared search city
- |— Current residential zip code
- |— Geographic mobility (local or national search)

PROJECT PROFILE

- |— Declared segment: primary residence / investment
- |— Target unit type: 1BR, 2BR, 3BR...
- |— Indicative budget (if provided)
- |— First-time buyer: yes / no / unknown
- |— Family situation (if provided)

DIGITAL BEHAVIOR

- |— Developments viewed (list + frequency + last visit)
- |— Time spent on financial calculators
- |— Documents downloaded (brochures, floor plans, FHA guides...)
- |— Previous emails: opens, clicks, on which content
- |— Acquisition source (paid search, real estate expo, referral...)

RELATIONAL STATUS

- |— Stage in journey (cold lead / warm lead / reserved / signed)
- |— Date of last contact
- |— History of interactions with sales advisors
- |— Any incidents (delivery delay, dispute...)

FINANCIAL DATA (partial)

- |— Income (if estimate was run)
- |— Estimated down payment (if provided)
- |— FHA eligibility calculated (if full estimate completed)

Example 1: Chatbot

The scenario is as follows: a user arrives on the website chatbot, unknown, and types *"I'm looking for an apartment in Denver to invest in."*

Step 1: Incoming Data and Initial Qualification

The incoming message is the only available signal. The agent has no access to any history, no CDP profile, no behavioral data. Two words ("invest" and "Denver") are the only elements available to trigger orchestration.

Step 2: Activation and Population of Each Context Layer

1. Brand & positioning

The developer operates across two deliberate and distinct segments: assisted homeownership and premium residential. The agent's tone must therefore be **bifurcated** by segment in the prompt. With a first-time buyer, the register is reassuring, pedagogical, supportive. The language of *"your home"* and *"your project"* is used; complexity is reduced. With a high-end investor, the register becomes precise, factual, performance- and wealth-

oriented: the language of "rental yield," "10-year appreciation," and "tax optimization" is used.

→ **Fallback:** if the client segment is unknown, the agent adopts a neutral tone and asks a qualifying question at the first exchange. "Is your project for your own use or for investment?" This is not a casual question; it is a context activation question.

The system prompt integrates these two distinct personas, and it is the customer data (profile, segment) that tells the orchestrator which one to activate.

2. Regulatory

The industry is particularly constrained. The agent can never guarantee the approval of a loan, an assistance program, or a tax benefit. It cannot assert a future resale price. It must systematically note that tax programs (FHA, cost segregation, 1031 exchanges, etc.) are subject to conditions and may change. Delivery timelines are estimates, the agent never commits to firm dates.

These rules are non-negotiable and hardcoded in the system prompt, with standard formulations for each sensitive situation.

3. Internal compliance

The property developer has its own rules: never disparage a named competitor, do not communicate on developments not yet permitted by the building permit, do not promise neighborhood improvements whose realization depends on third parties. There may also be rules around the handling of incoming leads. For example: every qualified contact must be passed to a human advisor within 24 hours; the agent does not close a sale on its own.

4. General & local market

This is where real estate development becomes particularly interesting. Two customers buying the same unit type in Austin and in Denver are operating in radically different markets in terms of supply pressure, price per square foot, rental dynamics, and competitive landscape.

→ **Interaction with customer data:** the customer's city (or zip code), or alternatively the city of the development they are interested in, is the lookup key. It triggers the retrieval of a structured market context: City = "Denver".

```
City: Denver
Market tension: strong, short days-on-market
Average new construction price: $420/sq ft
Rental dynamics: avg. gross yield ~4.4% before taxes
Competitive pressure: 3 active developers on mid-range segment
Recent signal: modest rate easing, renewed first-time buyer demand
```

The agent integrates this context to calibrate its messaging. In a tight market like Denver, it can create urgency ("availability on this type of development moves quickly"). In a softer market, it favors quality of service and project confidence.

→ **Fallback:** if the city is not provided, the agent uses the national market context as a base and steers the conversation quickly toward obtaining a location. "What area are you focused on?": a question that is both natural and functionally critical.

5. Business & offer

The agent queries the available developments in real time: location, available unit types, pricing, estimated delivery date, eligibility for assistance programs (FHA, down payment assistance, reduced rates...), unit status (available, reserved, under contract).

→ **Interaction with customer data:** the customer's financial profile (declared income, family situation, first-time buyer or not) is the filter key. The orchestrator cross-references the profile with the development database to surface only eligible properties. A first-time buyer in an FHA-eligible area with two children does not see the same inventory as a cash investor looking for cost segregation benefits.

→ **Fallback:** if the financial profile is incomplete, the agent presents a broad selection while explicitly noting the eligibility conditions for each development and invites the customer to clarify their situation to refine.

6. Channel & moment

In real estate, the channel matters particularly because the decision cycle is long and emotionally charged. An SMS is appropriate for a short re-engagement ("*Your saved development has a new unit available*"). An email allows rich content with floor plans and financial projections. A real-time chat is the right channel for qualifying and reassuring, not for closing.

→ **Interaction with customer data:** the customer's stated contact preferences, their past behavior (they open emails but never respond to texts), and their position in the journey (early exploration vs. approaching decision) determine which channel and format the agent uses.

→ **Fallback:** without a stated preference, the agent uses the incoming channel as a signal. If the customer initiated via chat, the agent stays on chat but offers to follow up by email for documents.

7. Agent objective & mission

In real estate, the mission varies significantly depending on the moment in the journey:

- **Exploration phase:** the objective is to qualify without pressuring, understand the project, build trust. Price is not raised before the project is understood.
- **Selection phase:** the objective is to refine, present 2-3 relevant options, and rule out alternatives without frustrating.
- **Decision phase:** the objective is to remove the last barriers, facilitate the financial estimate, prepare the handoff to a human advisor.
- **Post-reservation phase:** the objective shifts to reassurance and follow-up. Post-purchase anxiety is real in new construction with long delivery timelines.

→ **Interaction with customer data:** the journey stage (first visit, Nth interaction, unit reserved, contract signed) is the key that activates the right mission. The orchestrator reads this status and injects the corresponding task prompt.

→ **Fallback:** if the stage is unknown (new contact with no history), the agent always starts in exploration mode, with an open posture and broad project questions.

8. Relational & intentional

This is the richest layer in real estate, because profiles are highly contrasted and the emotional stakes are high.

A first-time buyer who returns to the same development page three times in a week sends a strong signal of intent, but also of hesitation, probably a financial obstacle or a fear of commitment. The agent does not push commercially and rather reassures and offers a concrete estimate.

An investor comparing three different cities in two days is in the market-selection phase. The agent provides comparative yield data rather than a description of the property.

A customer who experienced an issue on a previous development (delivery delay) approaches a new interaction with legitimate wariness. The agent acknowledges this explicitly before anything else.

Sample context summary:

```
Profile: Investor, 2 properties already purchased outside this developer
Current situation: Comparing Denver and Austin, 4 visits in 10 days
Behavioral signal: Long time spent on yield calculators
Historical incident: Delivery delay on development X in 2021
Probable state: Interested but wary about timelines
Probable intent: Decision in 2-4 weeks, looking for factual reassurance
```

→ **Fallback:** without history, the agent builds this profile progressively during the interaction by asking natural questions with a dual function: engaging the conversation and feeding the context in real time.

9. Learning & performance

The long-term memory base stores patterns by segment and situation. For a first-time buyer hesitating on financing in an FHA-eligible area, it might surface:

- the three most frequent objections (insufficient down payment, uncertainty about the purchase process, concerns about assistance program eligibility),
- the formulations that have converted best on this segment,
- the optimal moment to offer a financial estimate in the conversation.

→ **Interaction with customer data:** the customer's full profile (segment, city, financial situation, journey stage) serves as a search vector. The system does not look for similar cases in the abstract, it looks for cases similar to this customer in this market context. An investor in Denver gains nothing from patterns drawn from interactions with first-time buyers in Cleveland: the differences are too fundamental to be useful.

→ **Fallback:** if the profile is too sparse for a precise vector search, the agent uses learnings at the broad segment level (investor / assisted homeownership / premium residential) rather than at the granular level. Less precise, but always better than no learning at all.

Step 3: The Assembled Flow: Investor Visitor Case

```

INCOMING DATA: message + no history
↓
ORCHESTRATOR
├─ "invest" → activates Investor persona
├─ "Denver" → Denver market lookup
│   → strong tension, ~4.4% yield, rates easing
├─ Financial profile: unknown → broad presentation, light filtering
├─ Journey stage: unknown → mission: Exploration
├─ History: none → vector search on Investor Denver segment
│   → surfaces: frequent objection on depreciation benefits,
│   best hook = market data before product
├─ Channel: incoming chat → conversational format, concise
↓
ASSEMBLED PROMPT
→ Investor persona + contextualized Denver market
  + exploration mission + segment best practices
↓
AGENT: responds with Denver market data,
       asks a question about investment horizon
       before mentioning a specific development
  
```

Customer data is essentially zero at entry. It is the context, well-architected and well-orchestrated, that already enables a relevant, differentiated response aligned with segment best practices.

Step 4: Documented Fallback Decisions

In this scenario, the visitor is completely unknown. Almost all context layers tied to the customer profile therefore operate on fallback. The following documents the decisions made by the orchestrator in the absence of individual data, and their impact on response quality.

Missing data	Fallback applied	Impact
Customer segment unknown	Inferred from the word "invest." Investor persona activated. Qualifying question planned if ambiguous.	Low: strong linguistic signal
Financial profile unknown	Broad development selection, light filtering. Budget qualified through the interaction.	Medium: compensated by interaction
Journey stage unknown	Exploration mission activated by default. Agent qualifies before proposing.	Low: coherent posture
Interaction history: none	Learning at broad segment level (Investor, Denver). Generic but relevant patterns.	Medium: acceptable at start of journey
Channel preferences unknown	Incoming channel (chat) used as signal. Conversational format, concise, no commitment.	None: the channel is itself informative

Example 2: Relationship Marketing Email for an Investor

Situation: The developer is launching an email campaign on its new Denver development. Michael Chen, 47, is in the CDP. He has visited the development page twice, downloaded the brochure. Declared segment: investor. No estimate run. Last email opened three weeks ago, no clicks.

What the Channel Change Structurally Alters

The shift from chatbot to email fundamentally changes the nature of the problem: no back-and-forth is possible, no real-time inference, no qualifying questions. Everything must be resolved before generation. Context must compensate for the absence of interaction.

While in chat mode the agent can ask questions to fill context gaps, in email mode every gap must be resolved before generation, either through CDP data, explicit fallback rules, or a decision not to send. The orchestrator bears a considerably heavier load, and the quality of the final prompt depends entirely on the quality of this preparatory work.

Step 1: The Orchestrator Reads the CDP and Qualifies Each Dimension

```
CDP READ: Michael Chen
├─ Development city: Denver ✓
├─ Segment: Investor ✓
├─ Unit type: not provided X
├─ Budget: not provided X
├─ First-time buyer: not applicable (investor)
├─ Behavior: 2 visits to Denver development,
|           brochure downloaded, 0 estimates run
├─ Previous email: opened, 0 clicks → interest low or barrier unresolved
├─ Journey stage: warm lead, advanced exploration
├─ Incidents: none
├─ Financial data: none
```

Step 2: The Orchestrator Activates and Populates Each Context Layer

1. Brand & positioning

Investor persona activated. Tone: factual, performance- and wealth-oriented, no emotional register.

2. Regulatory

Injected as standard. Mandatory disclosures on tax strategies, no yield guarantee.

3. Internal compliance

Rule activated: no competitor mention, no per-square-foot pricing without "from" qualification.

4. Local market

Key: Denver → Lookup triggered:

Denver market / rental investment

- └ Rental tension: strong, vacancy <3.2%
- └ Average new construction gross yield: 4.1-4.8%
- └ Trend: modest rate easing over 60 days,
| renewed investor demand observed
- └ Competitive signal: competitor launching a comparable
| development in the same submarket Q1 2025
- └ Seasonality: favorable period (Q1, project reactivation)

→ This context will be used to create relevance and a modest sense of urgency without manufacturing it.

5. Business & offer

Key: Denver development + investor segment + available unit types

Denver development – investor inventory

- └ 1BR available: 4 units, from \$385,000
- └ 2BR available: 2 units, from \$495,000
- └ Estimated delivery: Q3 2026
- └ Cost segregation eligibility: yes on select units
- └ Estimated rental yield: 4.4% gross (1BR), 4.1% (2BR)
- └ Alert: 1 one-bedroom went under contract last week → inventory tightening

→ **Fallback:** unit type not recorded in CDP → present 1BR by default (best yield, most liquid, lower entry point → standard investor logic). Mention 2BR as a secondary option.

6. Channel & moment

Channel: email → Preferences: none declared → Past behavior: email opened, zero clicks → the subject line works, the content hasn't converted yet

Email channel rules activated:

- └ Length: medium (300-400 words), not long-form
- └ Structure: market hook → development benefit →
| single clear CTA
- └ Subject line: must be factual and numeric (investor profile)
- └ Send timing: Tuesday or Thursday morning (segment best practice)
- └ Behavioral diagnosis: previous email didn't convert →
| hypothesis: insufficient concrete yield data.
| This email must address that gap.

7. Objective & mission

Stage: warm lead, demonstrated interest (2 visits + brochure) without conversion → Mission activated: re-engagement toward action, not re-introducing the development

Step 3: The Assembled Prompt

SYSTEM PROMPT

[Investor persona + regulatory rules + internal compliance]

INJECTED CONTEXT

You are drafting a re-engagement email for Michael Chen, 47, a real estate investor.

MARKET CONTEXT:

Denver shows strong rental tension (vacancy <3.2%) and an average new construction gross yield of 4.1-4.8%. Rates have eased modestly over the past 60 days, which is reigniting investor demand.

DEVELOPMENT CONTEXT:

1BR units available from \$385,000, estimated yield 4.4% gross, delivery Q3 2026, cost segregation eligible. Inventory tightening (1 unit went under contract recently).

CLIENT CONTEXT:

Has visited the development twice, downloaded the brochure, never run an estimate. Previous email opened, no clicks. Probable barrier: insufficient concrete financial data. Do not repeat the brochure content.

MISSION:

Trigger an estimate consultation or a call. Bring new, specific, numeric information. Factual tone, not emotional.

LEARNING:

Embed an indicative projection in the email rather than redirecting to an external tool. Numeric, geolocated subject line. Address cost segregation proactively – it is the most frequent barrier at this stage.

FORMAT:

Email, 300-400 words, single CTA: "View the projection for this development."

Step 4: Documented Fallback Decisions

Missing data	Fallback applied	Impact on quality
Desired unit type	1BR by default (investor logic)	Low: consistent with segment
Budget	Not mentioned; price presented as entry point	Medium: less financially personalized
Financial estimate	Indicative projection built from market averages	Acceptable: sufficient behavioral signal
Email preferences	Segment best practices applied	Low: behavioral data compensates
Tax profile	Cost segregation mentioned with conditions, not personalized	Medium: regulatorily safe

Results

The system produces an individualized, contextualized email that addresses the real question the customer has not yet asked.

Denver: 4.4% estimated gross yield — and a tax angle the brochure didn't cover

From Sarah Mitchell — Investment Advisory <s.mitchell@[developer].com>
To Michael Chen <m.chen@[email].com>
Date Thursday, February 27, 2026, 8:42 AM

Michael,

You've looked at our Denver development twice and downloaded the brochure. I'm reaching out because there are aspects the brochure doesn't cover — and they tend to be the deciding factors at this stage of an investment decision.

DENVER RENTAL MARKET — CURRENT DATA

<p>VACANCY RATE < 3.2 % Strong demand, tight supply</p> <p>INTEREST RATE TREND Easing 60-day softening, renewed investor demand</p>	<p>AVG. GROSS YIELD — NEW CONSTRUCTION 4.1 – 4.8 % 1BR / 2BR residential segment</p> <p>AVAILABLE UNITS 6 units remaining 1 one-bedroom under contract last week</p>
--	--

INDICATIVE PROJECTION — 1BR, \$385,000, RINO / DENVER

<p>Estimated Monthly Rent \$1,850 Local 1BR rental market</p>	<p>Estimated Gross Yield 4.4 % Before taxes & expenses</p>	<p>Estimated Delivery Q3 2026 Fast path to income</p>
--	---	--

Projection based on current Denver rental market averages. Not a guarantee of returns. Final results depend on your financing terms and tax situation.

ON THE TAX SIDE — WHAT COST SEGREGATION CHANGES IN PRACTICE

Several units in this development qualify for accelerated depreciation through cost segregation — a strategy that can significantly improve after-tax returns in the early years of ownership. Eligibility and benefit magnitude depend on your tax bracket and holding structure. A personalized projection can shift the net yield materially depending on your marginal rate. I can prepare that analysis if you share a few details.

The timing is relevant: the rate easing underway since January is shifting financing math back in investors' favor. On this development, one 1BR went under contract last week — four remain as of today.

View the Personalized Projection for This Development

Response within 24 hours - No commitment required

I'm available if you'd like to go deeper on any specific angle — financing, tax structure, or delivery timeline.

Best regards,

Sarah Mitchell
 Head of Investment Advisory
 [Developer] — Investor Relations
 +1 (720) XXX-XXXX

Projected yields are indicative estimates based on market data and do not constitute a guarantee of returns. Tax strategies mentioned (cost segregation, accelerated depreciation) depend on individual eligibility and applicable tax law. Consult a licensed tax advisor for personalized analysis. - You are receiving this email because you expressed interest in our developments. [Unsubscribe](#)

Writing Choices

A few choices generated by the LLM from the prompt are worth unpacking:

The subject line does not say "*Our Denver Development*": it delivers new information (4.4%) and previews content the brochure did not cover. This is the learning for this segment: the numeric, geolocated subject line outperforms.

The opening explicitly acknowledges that the customer has already done the work of getting informed. It does not re-sell the development; it justifies why this email exists and what it adds.

The projection is in the body of the email, not linked to an external tool. This is the converting pattern for this profile according to long-term memory.

The cost segregation block proactively addresses the tax objection, identified as the #1 barrier at this stage, without waiting for Michael to raise it. In email, you cannot wait for the question.

The traceability panel below documents for each editorial choice which context layer produced it, using a color code: CDP (blue), dynamic lookup (green), learning (orange), fallback (purple). This information is logged in production to audit and improve the system.

INJECTED CONTEXT – TRACEABILITY LOG

Subject line	LEARNED	Numeric + geolocated format. Top-performing pattern for the investor segment according to long-term memory base.
Opening hook	CDP	Explicit reference to 2 visits + downloaded brochure. Avoids repeating content the prospect already has.
Market data	LOOKUP DENVER	Vacancy <3.2%, gross yield 4.1–4.8%, 60-day rate easing. Retrieved via local market intelligence tool.
Projection	LOOKUP DEVELOPMENT FALLBACK	1BR selected by default (best yield, most liquid, lowest entry point – standard investor logic). Income unknown → projection uses market averages.
Cost segregation block	LEARNED	Tax friction identified as the #1 objection at this stage for this segment. Addressed proactively without waiting for the question.
Inventory signal	LOOKUP DEVELOPMENT	"1 one-bedroom under contract last week" – real-time signal, not manufactured urgency.
CTA	LEARNED	"View the projection" outperforms "Learn more" and "Schedule a call" for this profile according to memory base.
Overall tone	CDP	Investor persona activated. Factual, data-driven, no emotional register. Words like "dream home" or "your place" deliberately absent.
Legal disclosures	REGULATORY	Non-guaranteed returns disclaimer + tax strategy eligibility conditions + compliance notice auto-injected.
Sender identity		Named human signature (Sarah Mitchell) – investor segment best practice: no generic sender address.

What This Case Reveals

In email, the orchestrator makes irreversible decisions: there will be no correction possible mid-interaction. The quality of the system therefore turns on three things:

- The richness of the CDP to minimize fallbacks
- The precision of the fallback rules so that default decisions remain relevant
- The quality of the long-term memory so that learning compensates for what individual data cannot provide

Example 3: Relationship Marketing Email for a Family Purchase

Step 1: The Orchestrator Reads the CDP and Qualifies Each Dimension

```

CDP READ: Alice Morgan
├─ Development city: Denver ✓
├─ Segment: not explicitly declared X
|   → strong behavioral signal: 2BR/3BR, upper floors,
|   budget >$500K → inference: primary residence
|   or premium primary residence
├─ Family profile: couple, 2 children, 35 years old ✓
├─ Professional situation: dual high-income earners ✓
├─ Budget: >$500K ✓
├─ First-time buyer: not recorded X
|   → dual high income + budget >$500K → very likely
|   already a homeowner or established purchasing capacity
├─ FHA eligibility: ineligible ✓ (income too high)
├─ Behavior: brochure + 2BR and 3BR upper-floor floor plans
|   → signal of concrete projection in the unit,
|   not generic browsing
├─ Previous email: not recorded → no email history
├─ Journey stage: hot lead, advanced selection phase
├─ Incidents: none

```

Step 2: Activation and Population of Each Context Layer

1. Brand & positioning

Premium primary residence segment activated (budget >\$500K + family profile + upper floors). Tone: warm but precise, oriented around quality of life, quality of build, family project. The language is "your family," "your daily life", not "yield." Aspirational register but grounded in the concrete: Alice is a senior professional; she reads through hollow marketing copy.

2. Regulatory

Injected as standard. No guarantee on delivery timelines, no commitment on resale values, mandatory disclosures on new construction purchase agreements. FHA not applicable, no mention to be made on assistance programs (risk of confusion).

3. Internal compliance

No competitive comparison. Pricing mentioned with "from" qualifier. Views and amenities on upper floors are presented as development features, not as definitive guarantees before closing.

4. Local market: Denver

Key: Denver + premium primary residence segment

Denver market / premium primary residence segment

- └─ Dynamic: active market on >\$500K segment,
genuine scarcity of quality new construction on upper floors
- └─ Premium new construction pricing: \$450-\$550/sq ft
- └─ Trend: modest rate easing over 60 days,
reactivating family purchase projects
- └─ Signal: few new developments available on this segment
in Denver right now – real scarcity
- └─ Local family context: neighborhood sought after for
quality of life, proximity to schools, urban vibrancy

5. Business & offer

Key: Denver development + 2BR/3BR + upper floors + budget >\$500K

Denver development – premium primary residence inventory

- └─ 2BR upper floors: 2 units available,
from \$495,000, 900 sq ft, open views
- └─ 3BR upper floors: 1 unit available,
from \$625,000, 1,180 sq ft, dual exposure
- └─ Amenities: premium finishes, private terrace,
covered parking included, storage
- └─ Estimated delivery: Q3 2026
- └─ FHA: not applicable (ineligible) → no mention
- └─ Inventory alert: the 3BR is the sole remaining unit on this floor

- **Orchestration decision:** lead with the 3BR (single unit, compatible budget, family of 4 profile). Mention 2BR as an alternative. The scarcity of the 3BR is a real signal, not a pressure technique.

6. Channel & moment

Channel: email → Email history: none → no actionable behavioral data → Profile: high-income professionals, limited available time

Email channel rules activated:

- └─ Length: concise and dense, not long-form
(high-income profile: fast reading, goes to the point)
- └─ Structure: daily-life projection → concrete property data
→ scarcity → CTA
- └─ Subject line: anchored in life project, not product
- └─ Send timing: Thursday morning (high-income segment best practice)
- └─ No email history → no behavioral diagnosis available,
segment best practices applied by default

- **Fallback (channel):** without email history, the performance patterns for the high-income family segment are applied. Short subject line, life-benefit oriented, no product title.

7. Objective & mission

Stage: hot lead, has downloaded precise floor plans of precise unit types on precise floors. This is no longer exploration, this is projection.

Mission activated: accelerate toward an in-person visit
 | or development walkthrough
 Objective: convert concrete interest into a qualified contact
 Tone of mission: create projection into daily life
 + real scarcity signal on the 3BR
 Prohibited: generic development pitch,
 repeating what the downloaded floor plans already show,
 mentioning FHA or any financial assistance
 (not applicable and potentially awkward)

8. Relational & intentional

Available signals:

- └ Both 2BR AND 3BR floor plans downloaded = she is comparing,
 | not yet decided on unit type → the 3BR warrants a case
- └ Upper floors exclusively = strong criterion,
 | probably views / light / quiet
- └ Budget >\$500K declared = capacity confirmed, no apparent
 | financial barrier
- └ Dual high income, 2 children, 35 = structuring life project,
 | joint decision, strong functional stakes
 | (bedrooms, spaces, family organization)
- └ FHA ineligible = no assistance constraints,
 | free choice over unit
- └ No incident = neutral relationship, first email interaction

9. Learning & performance

Search vector: primary residence + family + high-income + budget >\$500K + upper floors + advanced selection stage

Similar cases surfaced:

- └ Performing subject line: anchored in life project,
 | not product ("for your family" >> "3BR available")
- └ Converting content: concrete daily-life projection
 | (the kids' bedrooms, the terrace in the morning) combined
 | with precise factual data (sq ft, exposure, delivery)
- └ Primary lever at this stage: real inventory scarcity
 | (1 3BR remaining) + facilitating the joint decision
 | → offer a visit for both, not a solo call
- └ Frequent barrier: hesitation between 2BR and 3BR on budget
 | → argue the gap in terms of family use,
 | not in raw square footage
- └ CTA that performs: "Schedule a visit" >> "Learn more"
 | >> "Schedule a call" (too formal for this profile)

Step 3: The Assembled Prompt

SYSTEM PROMPT

[Premium primary residence persona + regulatory rules for new construction purchase + internal compliance]

INJECTED CONTEXT

CLIENT CONTEXT:

Alice Morgan, 35, dual high-income earners, 2 children. Budget >\$500K. FHA ineligible – no mention of assistance programs. Downloaded 2BR and 3BR upper-floor plans. Stage: active selection, comparing both unit types. Probable barrier: joint decision, 2BR vs. 3BR trade-off. Do not repeat the content of the already-downloaded plans.

MARKET CONTEXT:

Denver, genuine scarcity of premium new construction on this segment. Modest rate easing over 60 days.

DEVELOPMENT CONTEXT:

2BR upper floors: 2 units, from \$495,000, 900 sq ft, open views. 3BR upper floors: 1 unit remaining, \$625,000, 1,180 sq ft, dual exposure, private terrace, parking, storage. Delivery Q3 2026. Lead with the 3BR: last remaining unit, family of 4, compatible budget.

MISSION:

Help resolve the 2BR vs. 3BR trade-off by making family daily use concrete. Signal the real scarcity of the 3BR. Trigger an in-person visit for both.

LEARNING:

Concrete daily-life projection + precise factual data. Argue the 2BR/3BR gap in use terms, not square footage. CTA: "Schedule a visit for both of you."

FORMAT:

Email, 280-350 words, warm and precise tone, single CTA: "Schedule a Visit for Both of You."

Step 4: Documented Fallback Decisions

Missing data	Fallback applied	Impact
Segment not declared	Inferred as primary residence from behavioral signals (2BR/3BR, upper floors, budget, family profile)	Low: highly robust inference
First-time buyer unknown	Not mentioned in the email either way	None: not relevant for this profile
No email history	High-income family segment best practices applied	Medium: email optimized on profile type, not individual behavior
Precise income unknown	Budget >\$500K sufficient to qualify for 3BR. No financial estimate needed	None: the email is not financial

Missing data	Fallback applied	Impact
Joint decision-maker unknown	Implicitly addressed in CTA ("for both of you") and tone	Low: addresses the joint decision without naming it

The segment was not declared in the CDP. It was the orchestrator that inferred it from three combined signals: family profile, budget >\$500K, and the fact that she had exclusively downloaded upper-floor plans. A robust inference, but an inference nonetheless. If it were wrong (a family looking for a rental investment?), the entire email would be off-target. This is a strong argument for qualifying segment explicitly from the very first contact.

The FHA mention is deliberately absent, not because it is regulatorily prohibited, but because it would be tone-deaf. Reminding two high-income earners of their ineligibility for an assistance program, even implicitly, signals a failure of customer knowledge. Internal compliance context carries this type of subtle rule.

The barrier addressed is radically different: for Michael, it was the absence of concrete financial data. For Alice, it is the joint decision and the 2BR vs. 3BR trade-off. Same development, same city, same journey stage, two emails that share virtually no paragraph. This is precisely what properly architected context is designed to produce.

Results

The prompt produces the following email:

The top-floor 3BR — what the floor plans don't show for a family of four

From Claire Foster — Residential Advisory <c.foster@[developer].com>
To Alice Morgan <alice.morgan@[email].com>
Date Thursday, February 27, 2026, 8:47 AM

Alice,

You downloaded the floor plans for both the 2BR and 3BR upper-floor units in our Denver development. I'm writing to give you something the plans can't: how these two layouts actually live for a family of four, and why it may be worth making a call on the 3BR sooner rather than later.

WHAT THE 3BR CHANGES DAY TO DAY

With two kids, the question isn't really the extra 230 square feet — it's what those square feet enable. A bedroom for each child from day one. A home office or guest room that doesn't cannibalize the living area. And a private terrace where you can be outside without going downstairs. On the top floor, the dual exposure also transforms the quality of natural light through the day — something you feel when you come home in the evening and when you're actually living in the space on weekends.

2BR — Upper Floor
From **\$495,000** · 900 sq ft

- Open views, excellent light
- 2 bedrooms + living area
- Covered parking & storage included
- Estimated delivery Q3 2026

2 units available

RECOMMENDED FOR YOUR FAMILY

3BR — Top Floor
From **\$625,000** · 1,180 sq ft

- Dual exposure — light all day
- Private terrace
- 3 bedrooms — one for each child
- Covered parking & storage included
- Estimated delivery Q3 2026

Last unit available on this floor

The gap between the two is \$130,000. Spread over the life of a mortgage, and measured against what it enables for a growing family, it's often the decision people look back on and say it was worth every dollar.

On the Denver market right now: quality new construction at this level, on upper floors, is genuinely scarce. Mortgage rates have eased modestly over the past two months — a window that is bringing projects like yours back into favorable financing territory. The 3BR is the last unit on this floor; based on our sales history for comparable developments, this type of home rarely stays available for long.

Schedule a Visit for Both of You

Weekday and Saturday appointments available · Response within 24 hours

If you have questions before coming in — on delivery timing, finishes, or purchase process — I'm available by email or phone.

Best regards,

Claire Foster
Residential Advisor — Denver Developments
[Developer]
+1 (720) XXX-XXXX

Prices listed are indicative and may vary by unit and selected options. Delivery timelines are estimates and do not constitute a contractual commitment. New construction purchase subject to standard purchase agreement terms available upon request. - [unsubscribe](#)

The traceability panel below documents for each editorial choice which context layer produced it, using a color code: CDP (blue), dynamic lookup (green), learning (orange), fallback (purple). This information is logged in production to audit and improve the system.

INJECTED CONTEXT – TRACEABILITY LOG	
STABLE CONTEXTS	
Activated persona	INFERRED Premium primary residence. Segment not declared in the CDP – inferred from behavioral signals (2BR/3BR upper floors, budget >\$500K, family profile).
Tone	CDP Warm + precise. Oriented around quality of life and family project. Zero financial/yield vocabulary. Aspirational but factual register – high-income profiles read through hollow marketing copy.
Regulatory	Standard new-construction purchase disclosures in footer. Delivery timelines framed as "estimated." No mention of FHA or financial assistance programs – not applicable, and potentially tone-deaf for this profile.
DYNAMIC CONTEXTS	
Local market	LOOKUP DENVER Genuine scarcity of premium new construction on upper floors. 60-day rate easing. Premium primary residence segment (>\$500K) actively moving. Key: Denver + premium residential segment.
Development inventory	LOOKUP DEVELOPMENT 2BR: 2 units available, 900 sq ft, from \$495K. 3BR: 1 unit remaining, 1,180 sq ft, \$625K, private terrace, dual exposure. Delivery Q3 2026. 3BR inventory alert activated.
Decision: lead with 3BR	LOOKUP CDP Budget >\$500K compatible with 3BR. Family of 4 → 3 bedrooms argued on usage terms. Single unit remaining → real scarcity signal. 3BR lead decided by the orchestrator, not an editorial choice.
Channel & timing	FALLBACK No prior email history available → high-income family segment best practices applied. Subject line oriented toward life benefit, not product. Thursday morning send. Email kept dense and concise.
INFERRED & LEARNED CONTEXTS	
Mission	INFERRED Stage: active selection (specific floor plans downloaded). Mission: accelerate toward an in-person visit for both. No re-introduction of the development, no generic pitch.
Probable friction	INFERRED Joint decision + 2BR vs. 3BR trade-off. Not a financial barrier (budget confirmed). Email designed to facilitate that trade-off, not to convince her to buy.
Daily-life projection	LEARNED Top-performing pattern on family profiles: naming concrete everyday uses (a bedroom for each child, the terrace, the light) > listing raw square footage. Converts better than a spec sheet.
2BR vs. 3BR gap argument	LEARNED "\$130,000 spread over a mortgage" + family usage projection. Common friction on this segment: justify the gap in terms of everyday value, not absolute price.
CTA	LEARNED "Schedule a visit for both of you" outperforms "Schedule a call" and "Learn more" on family profiles in active selection. Implicitly includes the co-decision-maker without naming the dynamic.
Sender identity	Named female signature (Claire Foster) – consistent with profile and warm register. Title "Residential Advisor" vs. "Investment Advisory" used for Michael Chen.

Example 4: Relationship Marketing Email for First-Time Buyers

Step 1: The Orchestrator Reads the CDP and Qualifies Each Dimension

```

CDP READ: Tyler & Emma
├─ Development city: Denver ✓
├─ Segment: primary residence ✓
│   └─ → in-person visit + couple + 1BR + FHA eligible
│       └─ → first-time buyers highly probable
├─ Family profile: couple, unmarried, joint purchase ✓
│   └─ → specific legal implications (tenants in common /
│       joint tenancy) to note but not address in the email
├─ Financial situation: FHA eligible ✓
│   └─ → moderate to intermediate income confirmed
├─ Budget: not recorded X
│   └─ → FHA eligible + 1BR + Denver → probable range
│       $380,000-$430,000 (to confirm)
├─ First-time buyer: highly probable (FHA eligibility) ✓
├─ Digital behavior: interest in 1BR, in-person visit completed
│   └─ → advanced stage, real commitment – they made the trip
├─ Strong signal: "love the neighborhood" → emotional anchor
│   └─ → the neighborhood is already settled in their minds;
│       the question is whether to commit to this specific unit
├─ Previous email: not recorded → no history
├─ Journey stage: very hot lead, post-visit
│   └─ → critical moment: the post-visit window is where
│       people decide – or drift away
└─ Incidents: none

```

Step 2: Activation and Population of Each Context Layer

1. Brand & positioning

First-time buyer primary residence persona activated. This is the developer's most supportive and pedagogical register. The language is *"your first home,"* complexity is reduced, guidance is provided. Tyler and Emma are likely facing the largest financial commitment of their lives so far. The agent does not sell, it guides and reassures.

2. Regulatory

FHA: can be mentioned and explained, but always with the explicit caveat that final eligibility is subject to conditions verified by the lender. Never guarantee approval. Standard new construction purchase disclosures. The situation of an unmarried couple in a joint purchase has legal implications (title structure). The agent can note that a consultation with a real estate attorney is recommended but goes no further.

3. Internal compliance

Do not promise a definitive FHA estimate without verification of precise income. Do not address the legal question of the unmarried couple in an alarmist way. Signal, do not frighten. No competitive comparisons.

4. Local market: Denver

Key: Denver + first-time buyer segment + 1BR

Denver market / first-time buyer 1BR

- ├─ FHA-eligible area ✓
- ├─ Tension: strong on the affordable 1BR segment,
scarcity of new construction eligible for FHA
- ├─ New 1BR entry-level pricing: \$370,000-\$430,000
- ├─ Rate trend: modest easing over 60 days
 - materially improves monthly payment
for a first-time buyer with FHA financing
- ├─ Signal: the combination of FHA + rate easing
creates a real opportunity right now
- ├─ Neighborhood: dynamic, in-demand, well-connected
 - solid long-term value

5. Business & offer

Key: Denver development + 1BR + FHA eligibility + estimated budget \$380K-\$430K

Denver development – first-time buyer 1BR inventory

- ├─ 1BR available: 4 units, from \$385,000
 - FHA-compatible ✓
- ├─ Sizes: 640-760 sq ft
- ├─ Configurations: different floors and orientations
- ├─ Amenities: standard development, finishes included,
covered parking on select units
- ├─ Estimated delivery: Q3 2026
- ├─ FHA applicable: yes, subject to income conditions
 - to be confirmed with lender
- ├─ Inventory: 4 units available – reasonable availability,
but 1 went under contract (Michael Chen)
 - 3 one-bedrooms genuinely available today

→ **Orchestration decision:** do not push a single unit as with Alice. 3 units available means latitude. The urgency here is not inventory, it is the post-visit window. They loved the neighborhood; they made the trip. The risk is that they procrastinate or a competitor intercepts them.

6. Channel & moment

Channel: email → Email history: none → Profile: young couple, likely digitally active, sensitive to tone.

Email channel rules activated:

- |— Length: moderate, human and direct tone
| (not the dense investor register,
| not the premium register of Alice)
- |— Structure: post-visit emotional re-anchor →
| FHA financing made concrete →
| clear next step
- |— Subject line: should resonate with their visit experience,
| anchored in the neighborhood they loved
- |— Send timing: within 48 hours of the visit
| (critical memorization window)
- |— No history → first-time buyer post-visit best practices applied

7. Objective & mission

Stage: post-in-person visit. The most advanced stage possible before reservation. They came; they loved the neighborhood. The mission is not to convince them of the development but to translate the emotion of the visit into a concrete next step.

Mission activated: post-visit conversion

Objective: trigger an FHA financing estimate
or a second visit / model unit walkthrough

Tone of mission: warm, human, no pressure
"we're here with you" not "act fast"

Prohibited: generic development pitch,
artificial inventory alarm (3 units = no real urgency),
legal complexity around the unmarried couple
(signal, do not elaborate),
guaranteeing FHA eligibility without verification

8. Relational & intentional

Available signals:

- |— In-person visit completed = maximum commitment at this stage
| → they are not just curious; they are projecting themselves
- |— "Love the neighborhood" = strong emotional anchor
| → the neighborhood is already a settled decision;
| the question is: this development, this unit, now?
- |— Unmarried couple, joint purchase = joint decision
| → same dynamic as Alice but in a more emotional,
| less rational register
- |— FHA eligible = first-time buyers, probably
| first major joint financial decision
| → doubly stressful (first purchase + first purchase together)
- |— 1BR interest = life together, not yet family stage
| → do not project children into the messaging
- |— No incident = fresh, enthusiastic relationship

9. Learning & performance

Search vector: first-time buyer + couple + FHA + post-visit + 1BR + Denver

Similar cases surfaced:

- |— Best subject line: emotional re-anchor in the visit experience
| ("what you saw today" >> "1BR available")
- |— Converting content: FHA simplified as concrete benefit
| (monthly payment, reduced down payment)
| + emotional validation of the neighborhood
- |— Frequent barrier at this stage: "we'll think about it"
| → in reality: fear of administrative complexity
| and financing for a first purchase
| → solution: make the next step ultra simple and low-commitment
- |— Frequent error: sending a technical email after an emotional visit
| → register break → maintain the warmth of the visit in the email
- |— CTA that performs: "Run your FHA financing estimate"
| or "We'll call you to walk through the numbers"
| >> "Schedule a call" (too formal)
- |— Critical timing: email within 48 hours of visit
| → conversion rate ×2.3 vs. Day+7 email

Step 3: The Assembled Prompt

SYSTEM PROMPT

[First-time buyer primary residence persona +
new construction purchase regulatory rules +
conditional FHA framing + internal compliance]

INJECTED CONTEXT

CLIENT CONTEXT:

Tyler & Emma, unmarried couple, joint purchase.
FHA eligible – never guarantee approval;
always condition on lender verification.
Interest in 1BR; completed an in-person visit;
love the neighborhood. Stage: post-visit, very hot lead.
Probable first-time buyers – likely first major
joint financial decision.
Probable barrier: perceived complexity of financing
and commitment as an unmarried couple.
Do not address the legal question of the unmarried couple
in an alarmist way – simply note that a consultation
with a real estate attorney is recommended.
Do not project children into the messaging.

MARKET CONTEXT:

Denver FHA-eligible area.
Scarcity of new 1BR inventory eligible for FHA.
Rate easing over 60 days → improves monthly payment.

DEVELOPMENT CONTEXT:

1BR available: 3 units, from \$385,000,
640-760 sq ft, delivery Q3 2026. FHA applicable
subject to conditions. No artificial inventory urgency
– 3 units available but post-visit window is short.

MISSION:

Re-anchor the emotion of the visit.
De-dramatize FHA financing as a concrete benefit.
Make the next step ultra simple.
Trigger a financing estimate.

LEARNING:

Maintain the warm register of the visit.
Do not send a technical email after an emotional visit.
Simplify FHA: monthly payment, reduced down payment,
no jargon.
CTA: "Run your FHA financing estimate."

FORMAT:

Email, 250-300 words, human and direct tone,
warm without being saccharine.
Single CTA.

Step 4: Documented Fallback Decisions

Missing data	Fallback applied	Impact
Precise budget	Estimated range \$380K-\$430K on FHA + 1BR + Denver profile. No price emphasized in email. Estimate proposed to surface the budget.	Low: the estimate is the right tool to qualify
Precise income	FHA explicitly conditioned on lender verification. No specific projection in email body.	None: regulatorily sound and honest
Email history	First-time buyer post-visit best practices applied	Medium: compensated by the richness of behavioral signals
Partner's name	"both of you," inclusive without forcing an unknown name	None
Preferred floor / orientation	Not recorded. Not addressed in the email; reserved for next interaction.	Low: does not affect message relevance

The Tyler & Emma case reveals three things the previous two examples did not surface.

The in-person visit is the single most powerful CDP signal: it overrides almost all others in determining stage and mission. Someone who made the trip no longer needs to be convinced: they need to be guided toward a decision. The email does not sell; it removes obstacles.

FHA is treated as a de-dramatizing tool, not as an assistance program. For Michael it was absent; for Alice it would have been tone-deaf; for Tyler & Emma it is the central lever, translated into reduced down payment and accessible monthly payment, not administrative jargon. The regulatory context authorizes the mention; the learning context dictates how to frame it.

The unmarried couple is a CDP data point that triggers a specific regulatory context, the legal nudge. Neither Michael nor Alice has this block. It is present here not because of an editorial decision, but because the orchestrator cross-referenced "joint purchase, unmarried" with the internal compliance rules: signal, do not alarm, do not advise.

Results

The prompt produces the following email:

That neighborhood you liked — and what comes next

From Lucy Park — First-Time Buyer Guidance <l.park@[developer].com>
To Tyler <tyler@[email].com>
Date Thursday, February 27, 2026, 8:51 AM

Tyler,

You and Emma visited the development last week, and I heard you loved the RiNo neighborhood. I wanted to follow up — not to pitch the building again, but to show you what the next step actually looks like, and make it easier than it probably feels right now.

ON THE NEIGHBORHOOD — YOU'VE ALREADY MADE THE RIGHT CALL

RiNo has a decade-long track record of appreciation, walkability, and a sense of place that's harder to find in newer developments further out. The transit access, the restaurant corridor, the creative community — those aren't marketing talking points, they're what's driven consistent demand in this submarket. That part of the decision is done.

1BR — Denver RiNo Development

From 640 to 760 sq ft · Multiple floors and orientations
 Finishes included · Parking on select units
 Estimated delivery Q3 2026

From \$385,000
Starting price

3 units available

FHA FINANCING — WHAT IT MEANS IN PLAIN TERMS

<p>MINIMUM DOWN PAYMENT 3.5% <small>vs. 20% conventional — ~\$60K less upfront</small></p> <p>CREDIT SCORE FLOOR 580+ <small>For 3.5% down option</small></p>	<p>ESTIMATED MONTHLY (INDICATIVE) ~\$2,280 <small>Based on current rates + 3.5% down</small></p> <p>NEW CONSTRUCTION ELIGIBLE Yes <small>Subject to appraisal and lender approval</small></p>
---	---

FHA financing eligibility and terms are subject to verification by your lender. The figures above are illustrative estimates based on current market conditions. Final qualification depends on your income, credit profile, and debt-to-income ratio. We can connect you with a lending partner to run your actual numbers — at no cost and no obligation.

One thing worth a quick conversation with a real estate attorney: buying together as an unmarried couple involves a title and ownership structure decision — tenants in common vs. joint tenancy — that's worth sorting before you get to closing. It's straightforward, but it's better handled early. We can point you to someone if that's helpful.

Run Your FHA Financing Estimate

Takes about 5 minutes · No commitment · We'll walk you through it

The window right after a visit is usually when the picture is clearest. If you have questions — about the units, the process, or the financing — I'm here.

Warmly,

Lucy Park
 First-Time Buyer Guidance
 [Developer]
 +1 (720) XXX-XXXX

FHA financing eligibility and terms are subject to lender verification and applicable federal guidelines. Monthly payment estimates are illustrative and based on current market rates; actual payments will vary. New construction purchases are subject to standard purchase agreement terms. Ownership structure questions (joint tenancy, tenants in common) should be reviewed with a licensed real estate attorney. - Unsubscribe

The traceability panel below documents for each editorial choice which context layer produced it, using a color code: CDP (blue), dynamic lookup (green), learning (orange), fallback (purple). This information is logged in production to audit and improve the system.

INJECTED CONTEXT – TRACEABILITY LOG	
STABLE CONTEXTS	
Activated persona	CDP First-time buyer, primary residence. Confirmed via FHA eligibility. Warm, pedagogical, reassuring register – the most human-facing of the three personas.
Tone	LEARNED Human, direct, jargon-free. Pattern: never break the emotional register of the visit with a technical email. Identified as a frequent error on this profile.
FHA – regulatory framing	REGULATORY Mentioned and explained, never guaranteed. Always conditioned on "lender verification." Full disclaimer in footer and FHA block.
Unmarried couple	REGULATORY CDP Flagged via non-alarmist legal nudge. References title structure options without detail or alarm. Internal compliance: signal, don't counsel.
DYNAMIC CONTEXTS	
Local market	LOOKUP DENVER FHA-eligible new construction in RiNo is scarce. Rate easing over 60 days → direct monthly payment benefit for first-time buyers. Neighborhood long-term appreciation data integrated into validation block.
Development inventory	LOOKUP DEVELOPMENT 3 one-bedrooms available (4 minus 1 under contract – Michael Chen). From \$385K, 640–760 sq ft, delivery Q3 2026. FHA-compatible.
Decision: no inventory urgency	INFERRED 3 units = no artificial scarcity. Real urgency = post-visit window (48–72h). The lever is temporal and emotional, not inventory-based.
Channel & timing	LEARNED FALLBACK Email sent within 48h of visit (conversion rate ×2.3 vs. Day+7). No prior email history → first-time buyer post-visit best practices applied. Human tone maintained.
INFERRED & LEARNED CONTEXTS	
Mission	INFERRED Post-visit conversion. No re-selling needed (they already like it). Goal: make the next step feel simple and non-threatening.
Primary friction	INFERRED LEARNED Perceived complexity of financing a first purchase + commitment as an unmarried couple. Addressed via FHA simplification and soft legal nudge.
Neighborhood anchor	CDP "Love the neighborhood" = strong CDP signal. Explicitly validated in email to reinforce a decision already made. Avoid reopening a settled question.
FHA as concrete benefit	LEARNED Translated into reduced down payment + lower monthly cost, not regulatory language. Pattern that converts on first-time buyers vs. cold regulatory explanation.
No children projected	CDP Couple, no declared children. No family projection language in the copy. Common error on this profile: over-projecting an unsolicited future.
CTA	LEARNED "Run your FHA estimate" outperforms "Schedule a visit" on first-time buyers post-visit. Reduces perceived commitment. "No obligation" explicit under the button.
Sender identity	LEARNED Lucy Park, "First-Time Buyer Guidance" – reassuring functional title. Different from "Investment Advisory" (Michael) and "Residential Advisor" (Alice).
Budget undisclosed	FALLBACK No specific price emphasized. The financing estimate is the right tool to qualify budget – the email opens the door without assuming a precise number.

Comparison Across the Three Email Examples

The comparative table below places all three profiles side by side: same development, same city, same week. Zero paragraphs in common across the three emails.

THREE-WAY COMPARISON – SAME DEVELOPMENT, SAME WEEK, THREE CONTEXTS

DIMENSION	MICHAEL CHEN	ALICE MORGAN	TYLER & EMMA
Persona	Investor	Premium primary residence	First-time buyer
Segment – inferred or declared	Declared in CDP	Inferred (behavioral signals)	Inferred (FHA + visit + 1BR)
Dominant register	Factual, data-driven, wealth-oriented	Aspirational, anchored in family life	Warm, reassuring, pedagogical
Subject line	Yield figure + geolocation	Family lifestyle benefit	Emotional re-anchor post-visit
Primary lever	Market data + yield projection	3BR family use + real scarcity	Post-visit window + FHA simplified
Friction addressed	Missing concrete financial data	2BR vs. 3BR + joint decision	Financing complexity + unmarried couple
FHA / assistance programs	Absent (not applicable)	Absent (ineligible, would be tone-deaf)	Central block – simplified as benefit
Projection included	Rental yield in email body	None (friction is not financial)	FHA indicative in body, estimate via CTA
Urgency	Inventory (1 unit under contract)	Inventory (1 unit = last on floor)	Temporal (post-visit window)
CTA	View personalized projection	Schedule a visit for both of you	Run your FHA financing estimate
Sender / title	Investment Advisory	Residential Advisor	First-Time Buyer Guidance
Dominant color	Institutional navy	Premium steel blue	Accessible blue + reassuring green

Conclusion

The examples developed in this document make the case concretely: the quality of a CRM agent does not primarily depend on the language model chosen. It depends on how well the nine context layers are identified, fed, maintained, and orchestrated together. The same real estate development, addressed to three different customers in the same week, produces three emails with not a single paragraph in common. It is properly architected context, not superficial personalization, that makes this result possible.

Three structuring principles emerge from this analysis.

The first is the **separation between preparing and executing**: the orchestrator that assembles context and the agent that generates the response represent two distinct responsibilities and conflating them weakens the system.

The second is the **direct dependency between context quality and data quality**: a missing or poorly structured data point does not produce an incorrect response. It produces a fallback, meaning a generic response where a personalized one was achievable. Agentic systems surface the data gaps that humans used to compensate for through intuition.

The third is that **context is an asset that requires active maintenance**: stable layers must be governed and versioned; dynamic layers connected to reliable, current sources; learned layers fed through rigorous performance tracking.

The operational challenge that follows from this approach is as much organizational as technical. Defining and maintaining stable contexts requires cross-functional collaboration between brand, legal, and compliance teams. Feeding dynamic contexts requires reliable connectors to catalog systems, market intelligence tools, and channel management platforms. Building learned contexts requires a long-term memory infrastructure and a discipline of performance measurement. None of these responsibilities belong to the agent itself. They belong to the teams that deploy and operate it.

This document lays the foundation for a methodological approach to context. The natural next step is to derive from it an operational roadmap specific to each advertiser: which layers already exist, which must be built, what data is missing, what connectors are required. That mapping work is precisely where this framework leads.

Appendices

Sample Prompts for Example 1

Orchestrator prompt if the orchestrator is an LLM

Below is an example of a prompt as it would be submitted to an LLM orchestrator, that is, the prompt that asks a first model to build the context, before the surface agent speaks.

ORCHESTRATOR SYSTEM PROMPT

You are a context orchestrator for a real estate conversational agent. You do not speak to the customer. You prepare the context that will allow the surface agent to respond to them.

Your sole output is a structured context block, ready to be injected into the agent's prompt. You do not generate a response to the customer. You do not make unsupported assumptions. When information is missing, you apply the corresponding fallback rule and flag it explicitly.

AVAILABLE TOOLS

You may call the following tools:

- lookup_market(city, segment) → returns local market data for a given segment
- get_catalog(city, segment, max_budget, first_time_buyer) → returns available developments and their characteristics
- get_learnings(segment, city, journey_stage, behavior) → returns performance patterns from the long-term memory base
- get_channel_rules(channel) → returns format rules for the active channel

SITUATION

Channel: incoming chat

Message received: "I'm looking for an apartment in Denver to invest in"

CDP profile: none – unknown visitor, first interaction

YOUR MISSION

From the message and available tools, build the complete context block for the surface agent. Proceed in the following order:

Step 1: Extraction and qualification

Extract the available signals from the message. For each dimension, indicate whether the value is confirmed, inferred, or unknown.

Step 2: Tool calls

For each confirmed or inferred dimension with sufficient confidence, call the corresponding tool. For unknown dimensions, apply the fallback defined below and log it explicitly.

Step 3: Assembly

Produce the structured context block in the expected output format.

FALLBACK RULES

Segment unknown → infer from lexical signal if possible; otherwise activate neutral persona + mandatory qualifying question at end of response.

City unknown → use national market context as base; request location at the first exchange.

Budget unknown → broad catalog presentation, light filtering; do not mention pricing before qualification.

Journey stage unknown → activate Exploration mission by default.

History: none → vector search at broad segment level, not individual level.

EXPECTED OUTPUT FORMAT

ACTIVE_PERSONA: [value] – [confirmed / inferred / fallback]

MISSION: [value] – [reason]

MARKET_CONTEXT:

```
[lookup_market result or fallback note]
```

```
DEVELOPMENT_CONTEXT:
```

```
[get_catalog result or fallback note]
```

```
CHANNEL_CONTEXT:
```

```
[get_channel_rules result]
```

```
CLIENT_CONTEXT_SUMMARY:
```

```
[synthesis of available signals, probable state, probable intent]
```

```
LEARNINGS:
```

```
[get_learnings result]
```

```
FALLBACKS_APPLIED:
```

```
[list of dimensions on fallback and decision taken for each]
```

ORCHESTRATOR OUTPUT

Below is what the LLM orchestrator would produce in response to this prompt:

```
ACTIVE_PERSONA: Investor – inferred from lexical signal ("invest")
Confidence: high. No ambiguity in the message.
MISSION: Exploration – journey stage unknown, fallback applied

MARKET_CONTEXT:
[call lookup_market("Denver", "investment")]
→ Strong rental tension, vacancy < 3.2%
→ Average new construction gross yield: 4.1-4.8%
→ Modest rate easing over 60 days, renewed investor demand
→ Competitive signal: competitor development expected Q1 2025,
  same submarket
→ Seasonality: favorable period

DEVELOPMENT_CONTEXT:
[call get_catalog("Denver", "investor", max_budget=None, first_time_buyer=False)]
→ Light filtering applied (budget unknown)
→ 1BR: 4 units, from $385,000, estimated yield 4.4% gross
→ 2BR: 2 units, from $495,000, estimated yield 4.1% gross
→ Delivery Q3 2026 – cost segregation available on select units
→ Alert: 1 one-bedroom went under contract recently, inventory tightening

CHANNEL_CONTEXT:
[call get_channel_rules("chat")]
→ Conversational format, short responses (4-6 sentences)
→ Single question per message
→ No bullet points, direct tone

CLIENT_CONTEXT_SUMMARY:
Unknown visitor. Single signal: intent to invest in Denver.
Segment inferred, not confirmed. No history, no financial data.
Probable state: initial exploration, market discovery phase.
Probable intent: assess whether Denver is relevant before
going into the detail of a specific development.
Key note: do not go to product before qualifying investment
horizon and budget.

LEARNINGS:
[call get_learnings("investor", "Denver", "exploration", "first_contact")]
→ Performing hook: market data before product
→ Frequent objection at this stage: cost segregation / depreciation
  benefits (conditions, holding period)
→ Most effective qualifying question:
  "What's your investment horizon?" (less defensive than asking budget)
→ Optimal first response length: 4-6 sentences + 1 question

FALLBACKS_APPLIED:
- Budget: unknown → broad catalog, pricing not foregrounded
- Journey stage: unknown → Exploration mission activated
- History: none → broad segment learnings (Investor Denver)
- Financial profile: unknown → no FHA/depreciation eligibility filtering
```

This block is then concatenated with the agent's system prompt (brand, regulatory, compliance) to form the complete prompt passed to the surface agent.

Sample Prompts for Example 2

Below is the prompt submitted to the orchestrator, followed by its output, for the Michael Chen case.

ORCHESTRATOR SYSTEM PROMPT

You are a context orchestrator for a real estate marketing email generation agent. You do not write the email. You prepare the context that will allow the surface agent to write it.

Your sole output is a structured context block. You do not generate customer-facing content. In email mode, you have no ability to ask questions to fill context gaps – every missing dimension must be resolved by an explicit fallback decision, or by a decision not to send. There is no back-and-forth possible after generation.

AVAILABLE TOOLS

- get_cdp_profile(contact_id) → returns the complete contact profile from the CDP
- lookup_market(city, segment) → local market data
- get_catalog(city, segment, max_budget, first_time_buyer, unit_types) → filtered available developments
- get_learnings(segment, city, journey_stage, behavior) → performance patterns from long-term memory
- get_channel_rules(channel, email_behavior) → format rules + behavioral diagnosis

SITUATION

Channel: outbound email (campaign)

Development: Denver [name]

Contact: Michael Chen, contact_id = MC_4471

YOUR MISSION

Build the complete context block for the drafting agent.

Proceed in the following order:

Step 1 – CDP read

Call get_cdp_profile(MC_4471). For each dimension, indicate whether the value is confirmed, inferred, or unknown.

Step 2 – Fallback decisions

For each unknown or incomplete dimension, apply the fallback rule and document the decision. If a gap is too significant to produce a relevant email, indicate that it is better not to send.

Step 3 – Tool calls

Call tools in order, using confirmed or inferred CDP values as parameters.

Step 4 – Behavioral diagnosis

From the CDP data (previous email, digital behavior, stage), formulate a hypothesis about the probable barrier and instruct the agent accordingly.

Step 5 – Assembly

Produce the structured context block in the expected output format.

FALLBACK RULES

Unit type unknown → 1BR by default on investor profile (best yield, most liquid, lower entry point)

Budget unknown → do not mention a budget range; present price as an entry point

No estimate run → build an indicative projection from market averages,

```
present it in the email body
Email preferences unknown → apply investor segment best practices
(numeric + geolocated subject line; Tuesday or Thursday morning send)
Tax profile unknown → mention cost segregation with conditions,
do not personalize
Email opened without click → hypothesis: previous content too generic
or insufficiently numeric. This email must bring something new relative
to what the contact has already received or downloaded.
EXPECTED OUTPUT FORMAT
ACTIVE_PERSONA: [value] - [confirmed / inferred]
MISSION: [value] - [reason]
BEHAVIORAL_DIAGNOSIS: [synthesis + barrier hypothesis]

MARKET_CONTEXT: [lookup_market result]

DEVELOPMENT_CONTEXT: [get_catalog result + presentation decisions]

CHANNEL_CONTEXT: [get_channel_rules result + email diagnosis]

CLIENT_CONTEXT_SUMMARY: [CDP synthesis, probable state, probable barrier,
instruction to agent]

LEARNINGS: [get_learnings result]

FALLBACKS_APPLIED: [list of fallback dimensions, decision taken, estimated
impact]

DRAFTING_INSTRUCTIONS: [what the agent must do, what it must not do]
```

ORCHESTRATOR OUTPUT

Below is the output the orchestrator would produce:

ACTIVE_PERSONA: Investor – confirmed (declared CDP segment)
MISSION: Re-engagement toward action – warm lead stage confirmed,
demonstrated interest (2 visits + brochure) without conversion

BEHAVIORAL_DIAGNOSIS:

Michael did the work of getting informed: 2 development visits,
brochure downloaded. He has not committed. The previous email was
opened but generated zero clicks.

Barrier hypothesis: absence of concrete financial data about THIS
development. The brochure covers the property features, not the
numeric financial performance.

Instruction: do not repeat the brochure content. Bring something
he does not yet have – a yield projection embedded in the email body.

MARKET_CONTEXT:

[call lookup_market("Denver", "investment")]
→ Strong rental tension, vacancy < 3.2%
→ Average new construction gross yield: 4.1-4.8%
→ Modest rate easing over 60 days, renewed investor demand observed
→ Competitive signal: competitor development expected Q1 2025,
same submarket – decision window narrowing
→ Seasonality: favorable period (Q1)

DEVELOPMENT_CONTEXT:

[call get_catalog("Denver", "investor",
max_budget=None, first_time_buyer=False, unit_types=None)]
→ Unit type fallback applied: 1BR foregrounded (budget unknown,
standard investor logic)
→ 1BR: 4 units available, from \$385,000, estimated yield 4.4% gross,
delivery Q3 2026
→ 2BR: 2 units, from \$495,000, estimated yield 4.1%
(mention as secondary option)
→ Cost segregation eligibility: yes on select units (subject to conditions)
→ Inventory alert: 1 one-bedroom went under contract last week –
real signal, not manufactured

CHANNEL_CONTEXT:

[call get_channel_rules("email", behavior="opened_no_click")]
→ Length: 300-400 words
→ Structure: market hook → yield projection → development benefit →
single CTA
→ Subject line: factual and numeric (investor profile)
→ Send: Tuesday or Thursday morning (segment best practice)
→ Diagnosis: previous email opened, zero clicks.
Subject line works; content hasn't converted yet.
This email must be more data-dense than the previous one.

CLIENT_CONTEXT_SUMMARY:

Name: Michael Chen, 47
Segment: Confirmed investor
Behavior: 2 Denver development visits, brochure downloaded,
no estimate run
Previous email: opened, 0 clicks → real interest, barrier unresolved
Financial data: none
Incidents: none
Probable state: interested but blocked on the concrete financial
performance of the development
Probable intent: decision possible in the near term if the numbers

are compelling

Instruction to agent:

- Do not re-introduce the development (the client knows it)
- Embed a yield projection directly in the email, not via a link to an external tool
- Proactively address the cost segregation / depreciation objection (identified as #1 barrier on this profile at this stage)
- Create modest urgency on inventory (factual, not manufactured)
- Tone: factual, no persuasion, wealth-oriented

LEARNINGS:

- ```
[call get_learnings("investor", "Denver",
 "warm_lead", "brochure_without_estimate")]
```
- Outperforming subject line on this profile:  
numeric + geolocated ("Denver: 4.4% estimated gross yield")
  - Converting content: indicative projection embedded in the email body (no redirect to external tool)
  - Frequent barrier at this stage: uncertainty about depreciation benefits
  - CTA that performs: "View the projection for this development"  
>> "Learn more" >> "Schedule a call"
  - Timing: Day +21 after brochure download –  
Michael is at Day +18, optimal window

FALLBACKS\_APPLIED:

- Desired unit type: unknown
  - 1BR by default, standard investor logic
  - Impact: low, consistent with segment
- Budget: unknown
  - Price presented as entry point, not as a range
  - Impact: medium, less financially personalized
- Financial estimate: none run
  - Indicative projection built from market averages
  - Impact: acceptable, behavioral signal sufficient to justify
- Email preferences: none declared
  - Segment best practices applied
  - Impact: low, behavioral data compensates
- Tax profile: unknown
  - Cost segregation mentioned with conditions, not personalized
  - Impact: medium, regulatorily safe

DRAFTING\_INSTRUCTIONS:

TO DO:

- Open with a concrete Denver market data point (not the development)
- Embed an indicative 1BR projection in the body:  
\$385,000, 4.4% yield, rent and payment hypothesis
- Note the inventory signal (1 unit under contract) factually
- Address cost segregation with mandatory regulatory conditions
- Numeric, geolocated subject line
- Single CTA: "View the projection for this development"

NOT TO DO:

- Repeat the brochure content already downloaded
- Redirect to an external estimator (conversion barrier)
- Guarantee a yield
- Mention the competitor
- Manufacture artificial inventory pressure

The structural difference from Example 1 is visible in the length of the generated prompt. In the chatbot scenario, the orchestrator was working from almost nothing, inferring and applying fallbacks across nearly every dimension. Here, the CDP provides real signal on behavior, stage, and email history, which allows the orchestrator to produce a precise behavioral diagnosis and explicit drafting instructions. The quality of context is directly proportional to the richness of the input data.