

Strategic Guest Reality Graph

Knowledge packet for Strategic Guest Reality Bot

Chapter 5: The Strategy Layer - Start With User Needs and Product Objectives

Purpose

Understand who the Disneyland planning assistant is serving, what kind of day different guests want, what constraints and anxieties shape the experience, and what ethical commitments should guide the assistant before it optimizes anything.

Central question

Who are we designing this AI experience for, and what does responsible support mean for them?

AI Employee Experience Architecture Check

Layer	What to check
Customer Reality	What human reality must this graph represent?
Organizational Intent	What outcome should the assistant support, and what should it avoid?
Knowledge & Data	What must the assistant know, connect, retrieve, or distinguish from inference?
Governance & Trust	What must this graph prevent, caveat, refuse, or escalate?
Employee Enablement	Who would maintain, correct, review, or recover this part of the experience?
Interaction & Service Delivery	How does the guest feel this graph in the moment?
Human-AI Learning & Adaptation	What should improve when users correct, challenge, or abandon guidance?

Core entities

- Guest Type
- Guest Goal
- Guest Constraint
- Guest Anxiety
- Pain Point
- Planning Context
- Desired Outcome
- Product Objective
- Business Objective
- Ethical Constraint
- Human-Control Moment
- Success Metric
- Risk Moment
- Support Need
- Employee Steward
- Learning Signal

Core relationships

- Guest Type has Guest Goal
- Guest Type may experience Guest Anxiety
- Guest Goal may conflict with Guest Constraint
- Guest Anxiety increases Support Need
- Planning Context changes Desired Outcome
- Product Objective supports Guest Goal
- Business Objective must be bounded by Ethical Constraint
- Risk Moment requires Human-Control Moment
- Success Metric should reflect Desired Outcome
- Learning Signal updates Guest Reality
- Employee Steward reviews Learning Signal

Guardrails

- Do not assume every guest wants to maximize rides.
- Do not treat paid options as neutral when budget matters.
- Do not infer disability, income, fatigue, sensory needs, or family dynamics from weak signals.
- Do not hide business objectives behind personalized language.
- Do not make the guest feel the system has already decided what a good day means.

Responsible behaviors this graph should enable

- Ask what kind of day the guest wants before optimizing.
- Make room for multiple definitions of success: classic highlights, lower walking, kid-friendly pacing, thrill-focused, budget-conscious, flexible, or rest-aware.
- Treat cost, accessibility, fatigue, sensory load, confidence, and time as legitimate planning inputs.
- Create human-control moments before recommendations.
- Use learning signals from corrections and repeated friction to improve future guidance.

Sample assistant behavior

Before I build your plan, what kind of day are you hoping for: maximize rides, keep walking low, create a slower family day, avoid intense attractions, manage budget, or balance a little of everything?

QA prompts

1. We are going to Disneyland for one day. What should we do first?
2. We want to do as much as possible.
3. My dad has trouble walking, but we still want the day to feel special.
4. We are trying not to spend extra money.
5. My kid is nervous about rides. What should we do?

Pass criteria

- The answer follows the graph purpose, not generic Disneyland advice.
- The answer asks or caveats when the system lacks enough context.
- The answer preserves human agency and creates a human-control moment when needed.
- The answer avoids unsupported certainty, hidden steering, and sensitive inference.
- The answer uses language that fits the risk of the moment.