



Certified Chatbot Developer (CCHD)

Ten Sample Questions

These ten questions are written at the difficulty of the live exam. Six carry a figure, a data table, a configuration or a code listing; four are plain scenario questions, as on the live paper. Every one is a scenario a working chatbot developer would meet, and none of them appears in the live question bank.

The live exam is multistage adaptive: how you score on one stage selects the difficulty of the next, so no two candidates sit the same paper. It is a performance-based exam: each form can carry up to 5 to 12 performance-based questions alongside the multiple-choice questions, which means you build, sequence or complete something rather than only choosing from a list. Answers and full explanations follow each question here; the live exam does not disclose item-level answers.

Question 1

Domain 1: Chatbot Foundations

DATA

Sarah Miller maps each chatbot type to how it produces a reply, with one method missing.

Type	Reply method
Rule-based	Fixed scripted paths
Retrieval	(to choose)
Generative	Model writes new text
Menu	User taps buttons

Table 1.1 Chatbot types and reply methods.

How does a retrieval bot produce a reply?

- ☐ A Writes brand new text
- ☐ B Picks stored responses
- ☐ C Ignores the user input
- ☐ D Calls a human agent

Correct answer: B

A retrieval bot selects the best match from a set of stored responses rather than generating new wording. Generative bots write new text, rule-based bots follow scripted paths, and neither ignores input nor always escalates to a person.

Question 2

Domain 2: Natural Language Understanding

CONFIG

James Carter reviews the JSON result the NLU engine returns for one user message.

```
{
  "text": "book a flight to Denver",
  "intent": "BookFlight",
  "confidence": 0.94,
  "entities": [
    { "type": "city", "value": "Denver" }
  ]
}
```

Listing 2.1 An NLU result.

What does the entities array capture?

- ☐ A The overall user goal
- ☐ B The reply to send back
- ☐ C Key values in the text
- ☐ D The model version used

Correct answer: C

Entities capture the specific values pulled from the message, here the city Denver. The intent field holds the overall user goal, confidence scores the match, and none of these fields carries the outgoing reply.

Question 3

Domain 3: Dialog Design and Flows

CONFIG

Olivia Davis reviews the slot filling config for a pizza order flow.

```
order:
  slots:
    - name: size
      prompt: What size?
      required: true
    - name: topping
      prompt: Which topping?
      required: true
  confirm: true
```

Listing 3.1 A slot filling flow.

What does required true cause for a slot?

- ☐ A Bot has to fill it
- ☐ B Bot may skip the slot entry
- ☐ C Bot ends the chat now
- ☐ D Bot logs it silently

Correct answer: A

A required slot means the bot must keep prompting until it collects that value before moving on. It does not let the bot skip the slot, end the chat, or merely log the field without asking.

Question 4

Domain 4: Building and Integrations

CONFIG

Liam Anderson reviews the webhook payload the bot posts to an order service.

```
{
  "action": "createOrder",
  "channel": "web",
  "user": "guest-482",
  "items": [ "pizza", "soda" ]
}
```

Listing 4.1 A webhook payload.

What does this webhook let the bot do?

- ☐ A Train a new NLU model
- ☐ B Style the chat bubble
- ☐ C Store model weights only
- ☐ D Act in another system

Correct answer: D

A webhook lets the bot trigger an action in an external system, here creating an order in the order service. It does not train models, store weights, or change how the chat interface looks.

Question 5

Domain 5: LLM-Powered Chatbots

DATA

Ava Thompson compares ways to ground an LLM chatbot, with one description missing.

Approach	What it does
Fine-tuning	Retrains model weights
Retrieval augmented	(to choose)
Prompt only	Relies on base model
Guardrails	Filters the output

Table 5.1 Grounding approaches.

What does retrieval augmented generation do?

- ☐ A Edits the model weights
- ☐ B Adds fetched context
- ☐ C Removes all prompts
- ☐ D Turns off the model

Correct answer: B

Retrieval augmented generation fetches relevant documents at query time and adds them to the prompt as context. It leaves the model weights unchanged, still uses prompts, and keeps the model running.

Question 6

Domain 6: Testing, Analytics and Safety

CONFIG

David Brown reviews the safety config applied to the deployed bot.

```
safety:
  block_profanity: true
  pii_redaction: true
  max_tokens: 500
  escalate_to_human: true
```

Listing 6.1 A safety config.

What does pii redaction do?

- ☐ A Speeds up the replies
- ☐ B Adds more model tokens
- ☐ C Hides personal data
- ☐ D Logs every user message

Correct answer: C

PII redaction removes or masks personal data such as names or card numbers before it is stored or shown. It does not change reply speed, raise the token limit, or record every message.

Question 7

Domain 1: Chatbot Foundations

A team builds a bot that only follows fixed button menus and cannot handle a typed free-form question. What capability is missing?

- ☐ A More menu buttons here
- ☐ B A faster hosting server
- ☐ C A brighter chat theme
- ☐ D Free text input

Correct answer: D

Handling typed free-form questions requires natural language input rather than fixed menus alone. More buttons, faster hosting, or a brighter theme do not give the bot the ability to understand open text.

Question 8

Domain 2: Natural Language Understanding

A bot maps cancel my order and I want to stop this order to two separate intents, so replies feel inconsistent. What best fixes this?

- ☐ A Add more chat channels
- ☐ B Delete all the entities
- ☐ C Group into one intent
- ☐ D Raise confidence to one

Correct answer: C

Both phrasings express the same goal, so grouping their training examples under one intent gives a consistent reply. Adding channels, deleting entities, or forcing confidence to one does not address the duplicated intents.

Question 9 Domain 4: Building and Integrations

A bot needs to show a customer their live order status, but it only has scripted replies and no link to the order database. What does it need?

- ☐ A A backend link
- ☐ B More greeting phrases
- ☐ C A longer welcome text
- ☐ D A fresh avatar image

Correct answer: A

Live order status lives in a backend system, so the bot needs an integration that queries that data at request time. Extra greetings, longer welcome text, or a new avatar do not connect the bot to the order records.

Question 10 Domain 6: Testing, Analytics and Safety

After launch, a bot often fails to answer real questions, but the team never reviews chat logs or fallback rates. What practice is missing?

- ☐ A Change the bot color
- ☐ B Add more empty slots
- ☐ C Shorten every prompt
- ☐ D Review the chat logs

Correct answer: D

Reviewing chat logs and fallback rates reveals where real users are failing so the bot can be improved. Changing colors, adding slots, or shortening prompts does not surface those gaps in coverage.