DUMPSARENA

AWS Certified Alexa Skill Builder-Specialty

Amazon AWS AXS-C01

Version Demo

Total Demo Questions: 12

Total Premium Questions: 123

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Voice-First Design	22
Topic 2, Design Architecture	10
Topic 3, Development	60
Topic 4, Testing	16
Topic 5, Deployment and Operations	15
Total	123

QUESTION NO: 1

An Alexa Skill Builder is developing a complex skill to help users purchase items. The Builder has implemented the default AMAZON.HelpIntent, but when users ask for help they always reach the same unhelpful result.

How should the Builder enhance the experience to guide the users toward the end goal?

A. Check the state of the user conversation when AMAZON.HelpIntent is received to give contextual responses.

B. Implement AMAZON.FallbackIntent to give a response when an utterance is not understood.

C. Extend the default AMAZON.HelpIntent with more sample utterances.

D. Implement a custom intent users can request when they need help.

ANSWER: A

Explanation:

Check the state of the user conversation when AMAZON.HelpIntent is received to give contextual responses. AMAZON.HelpIntent is designed to let customers ask for assistance at any point in a skill, but Alexa does not automatically know what guidance is most useful for the skill's current purchase flow. The skill's handler should inspect its session attributes, dialog state, current intent, or other stored progression data and then return help specific to that point in the interaction. For example, when a customer is selecting an item, the response can explain how to choose or hear available items; when the customer is confirming a purchase, it can explain the available confirmation response. Contextual help reduces repetition, keeps the customer oriented, and gives an actionable next step toward completing the purchase. The handler can also use a reprompt that preserves the current step, allowing the customer to continue the transaction naturally after receiving assistance. Amazon identifies AMAZON.HelpIntent as the standard built-in intent for handling help requests, and skill guidance recommends tailoring responses to the customer's context. See Standard Built-In Intents and Manage Skill Sessions and Session Attributes.

QUESTION NO: 2

An Alexa Skill Builder supports a custom skill in both en-US and en-GB. The Builder adds a new intent and sample utterances to the en-US interaction model.

What must the Builder do for customers using the en-GB version of the skill to use the new intent?

A. Add the intent and sample utterances to the en-GB interaction model and build that locale.

B. Redeploy the en-US interaction model because Alexa automatically copies it to all supported locales.

C. Add the en-GB locale to the Lambda function environment variables.

D. Create a second Lambda function in the United Kingdom AWS Region.

ANSWER: A

Explanation:

The Builder must add the new intent and applicable sample utterances to the **en-GB interaction model** and build that locale. Each supported locale has its own interaction model, including its invocation name, intents, sample utterances, slot values, and dialog configuration. Updating en-US does not automatically update en-GB.

The Builder should use wording and slot values appropriate for the target locale, then test the en-GB model using the selected locale in the developer console. See the Develop Skills in Multiple Languages documentation.

QUESTION NO: 3

According to Amazon Alexa best practices, how should an Alexa Skill Builder prevent unintentional requests against a skill's backend when using AWS Lambda?

A. Ensure that the session ID provided by the request to Lambda is not already in use.

B. Rotate the Lambda ARN regularly to prevent others from using the service.

C. Retrieve the Application ID property from the request JSON and validate it against the Lambda environment variables.

D. Provide the Lambda trigger with the Application ID so that it validates on the ask trigger.

ANSWER: D

Explanation:

Provide the Lambda trigger with the Application ID so that it validates on the ask trigger.

is correct because associating the Alexa Skills Kit trigger with the skill's unique identifier constrains invocation permissions to that Alexa skill. When the trigger is configured, AWS Lambda creates or updates a resource-based policy that grants the Alexa Skills Kit service permission to invoke the function for the specified skill. This is the preferred infrastructure-level control: the backend is explicitly connected to the intended Alexa skill rather than being generally available to requests associated with other skills.

In current AWS terminology, this identifier is commonly called the Skill ID. It is supplied while configuring the Alexa Skills Kit trigger in the Lambda console or when creating the corresponding Lambda permission through infrastructure tooling. The resulting permission uses the Alexa service principal and scope for the designated skill, helping prevent accidental or unintended invocations of the function by a different Alexa skill. This configuration should be established as part of deploying the skill's Lambda endpoint and maintained if the skill or backend is changed. See the AWS documentation on using AWS Lambda with Alexa Skills Kit and Alexa's skill manifest and Skill ID documentation.

QUESTION NO: 4

Before beginning the development of a skill, an Alexa Skill Builder wants to hear how a statement will be pronounced by Amazon Alexa in another locale.

How should the Builder test the pronounciation of the statement?

A. Use the Alexa Simulator tab in the developer console to interact with the skill using the selected locale.

B. Create a developer account for the specific locale, and use a device to interact with the skill.

C. Use Voice & Tone tab in the developer console to play the statement using the selected locale.

D. Use the Manual JSON tab in the developer console to create a sample JSON response

ANSWER: C

Explanation:

Use Voice & Tone tab in the developer console to play the statement using the selected locale. The Voice & Tone tool is designed for previewing how Alexa renders speech without requiring a completed skill, interaction model, endpoint, or physical Alexa-enabled device. The builder can enter the proposed statement, select the target locale, and play the generated audio to assess pronunciation, pacing, and the locale-specific voice rendering. This is particularly useful before implementation because the same written text can be spoken differently across locales based on language rules, pronunciation conventions, and available Alexa voices. The preview also helps the builder decide whether the statement needs punctuation changes, phonetic guidance, or Speech Synthesis Markup Language (SSML) adjustments before it is included in a skill response. SSML supports controls such as language selection, pronunciation guidance, pauses, and prosody, which can be validated through speech preview as part of designing a natural customer experience. Amazon documents the developer-console testing tools at Test Your Skill and describes speech-output controls in the Alexa SSML Reference.

QUESTION NO: 5

An Alexa Smart Home Skill controls dimmable lights. During device discovery, the Builder needs to declare the Alexa capabilities that allow customers to turn the lights on or off and adjust brightness.

Which capability interfaces should the Builder include for each dimmable light endpoint? (Choose two.)

A. Alexa.PowerController

B. Alexa.BrightnessController

C. Alexa.TemperatureSensor

D. Alexa.PlaybackController

E. Alexa.LockController

ANSWER: A B

Explanation:

The `Alexa.PowerController` capability declares that the endpoint supports power operations such as `TurnOn` and `TurnOff`. Alexa uses this capability to route customer power commands to the smart home skill.

The `Alexa.BrightnessController` capability declares support for brightness adjustments, such as setting a light to a percentage value. The discovery response must accurately list the capabilities that the endpoint supports so Alexa can expose and route the appropriate customer commands. See the Alexa.PowerController interface reference and the Alexa.BrightnessController interface reference.

QUESTION NO: 6

An Alexa Skill Builder noticed that a large percentage of a food ordering skill's customers are not completing their transactions. The Builder needs to know what portion of the customers are leaving the skill by not responding, compared to the portion of customers who receive an error.

Which report inside the Analytics section of the developer console will provide this information?

A. Session Type Distribution

B. Unique Customers Per Intent

C. Average Session Per Customer

D. Failed Utterances Per Intent

ANSWER: A

Explanation:

Session Type Distribution is the appropriate Analytics report because it categorizes skill sessions by how they ended. This gives the Skill Builder a direct view of the relative share of sessions associated with user inactivity or abandonment—where the customer does not provide another response—versus sessions that end because an error occurs. For a food-ordering flow, this distinction is important: inactivity can indicate that a prompt is unclear, the conversation has become too long, or the customer needs more guidance, while error-ending sessions identify reliability or implementation issues that interrupt the transaction. Reviewing the distribution over a selected date range helps the builder quantify each outcome and prioritize improvements to the dialog design or error handling based on observed customer behavior. Amazon documents this report as part of the developer console's skill-usage analytics, which are intended to help developers understand customer engagement and session outcomes. See Measure Skill Usage and Analytics for Alexa Skills.

QUESTION NO: 7

An Alexa Skill Builder is developing a skill that tells users facts about their neighborhoods using their postal codes. When running the skill, users receive error messages from their Amazon Alexa devices. When the Builder checks the logs, error

messages are being sent when the device location API is used to get users' postal codes. What are the potential causes of this issue? (Choose two.)

A. The Alexa devices are registered in the incorrect postal code.

B. The AWS Lambda function is not registered to the skill.

C. The users have opted out of giving their location information.

D. The application ID of the skill does not match the expected application ID of the AWS Lambda function.

E. The skill is not registered to access information in the developer console.

ANSWER: C E

Explanation:

The users have opted out of giving their location information is a potential cause because Alexa treats device address data, including postal code, as customer-controlled information. A skill can retrieve this information only after the customer has granted the required permission for that skill. If the customer has not granted the permission, or later revokes it in the Alexa app, an address API request cannot return the postal code and can result in an authorization-related error response.

The skill is not registered to access information in the developer console is also a potential cause. Before requesting a device postal code through the Alexa Device Address API, the skill builder must enable the relevant device-address permission in the Alexa Skills Kit developer console. This permission configuration identifies the customer data the skill intends to request and allows Alexa to present the consent prompt. The skill should also obtain the API endpoint from the request context and use the request-scoped API access token when calling the endpoint. With both developer-console permission configuration and customer consent in place, the skill can request the device postal code for its neighborhood-fact experience.

See the Alexa Device Address API documentation and the customer permissions documentation.

QUESTION NO: 8

An Alexa Skill Builder is developing a skill that must send an initial welcome email to each new user.

What combination of features would the Builder use to satisfy this requirement? (Choose two.)

A. Alexa Settings API

B. Entity resolution

C. Customer Profile API

D. Device Address API

E. Data persistence

ANSWER: C E

Explanation:

The **Customer Profile API** provides the email address needed to deliver a welcome email. The skill must request the customer's permission to access email through the appropriate Alexa permission, and it should obtain the email address only after permission has been granted. Alexa documents that customer contact information, including email address, is accessed through the Customer Profile API and requires customer consent. The skill can then pass that address to an email delivery service or its backend email workflow.

Data persistence is needed to ensure that the message is initial rather than repeated. The skill should persist a per-user value, such as a welcome-email-sent flag, using the Alexa user ID as the logical user key. On a subsequent request, the skill reads that value before sending email. If no record exists, it sends the welcome message and stores the flag after successful processing. This approach reliably distinguishes a newly encountered user from one who has already received the message,

including across separate skill sessions and invocations. See the Alexa documentation for requesting customer contact information and persistent attributes and persistence.

QUESTION NO: 9

An Alexa Skill Builder is troubleshooting a custom skill backed by AWS Lambda. The Builder wants to trace calls from the Lambda function to a database and a third-party HTTP service to identify where request latency occurs.

Which actions should the Builder take? (Choose two.)

A. Enable active tracing for the Lambda function.

B. Grant the Lambda execution role permission to write trace data to AWS X-Ray.

C. Enable Amazon S3 versioning for the Lambda deployment package.

D. Configure an Amazon DynamoDB global secondary index for the function.

E. Increase the Alexa skill session timeout.

ANSWER: A B

Explanation:

The Builder should enable active tracing for the Lambda function. Active tracing causes Lambda to send trace data to AWS X-Ray for supported invocations, allowing the Builder to view execution segments and downstream calls.

The Lambda execution role must also allow the function to write trace data to AWS X-Ray. The role requires permissions for operations such as `xray:PutTraceSegments` and `xray:PutTelemetryRecords`. AWS provides the managed `AWSXRayDaemonWriteAccess` policy for this purpose.

See the AWS documentation for Using AWS X-Ray to trace AWS Lambda applications and Instrumenting Lambda functions with AWS X-Ray.

QUESTION NO: 10

An Alexa Skill Builder published a skill and now wants to update it.

Which of the following changes will make it necessary to have the skill re-certified? (Choose two.)

A. Adding a new intent in the interaction model

B. Changing the text of a skill response

C. Adding a sentence to the skill description

D. Changing the URL of the image on a skill card

E. Updating the backend to point to a new database

ANSWER: A C

Explanation:

Adding a new intent in the interaction model requires recertification because the interaction model defines the customer utterances, intents, slots, and dialog behavior through which users access a skill. A new intent changes the skill's certified voice experience and can introduce new functionality, so Amazon must review the updated model before the change is published to customers. Adding a sentence to the skill description also requires recertification because the description is part of the skill's public Store listing. Store metadata is reviewed to ensure that it accurately represents the skill and continues to meet Alexa certification requirements. Amazon's guidance for updating published skills distinguishes changes to the interaction model and skill Store details from operational changes made solely within the skill's hosted implementation.

Submit these updates through the Alexa Developer Console and complete the certification workflow before publishing them. See the Amazon documentation for updating a published skill and the Alexa Skills certification requirements.

QUESTION NO: 11

An Alexa Skill Builder supports a custom skill in en-US and de-DE. The Builder wants customers in each locale to receive recognition and responses appropriate for their language.

Which actions should the Builder take? (Choose two.)

A. Create and build a separate interaction model for each supported locale.

B. Use the locale in the incoming request to select localized response content.

C. Use the en-US interaction model automatically for every supported locale.

D. Create a separate AWS Lambda function for each language.

E. Configure all localized utterances as synonyms in one custom slot type.

ANSWER: A B

Explanation:

The Builder should create and build an interaction model for each supported locale. Interaction models are locale-specific and contain language-dependent resources such as invocation names, intents, sample utterances, slot values, prompts, and dialog configuration.

The endpoint should also inspect the locale in the incoming request and return localized output appropriate for that locale. The locale is available in the request object and identifies the language and region in which Alexa processed the request. The skill can use it to select the correct response strings, SSML, and locale-specific backend behavior.

See the Develop Skills in Multiple Languages documentation and the Alexa request and response JSON reference.

QUESTION NO: 12

An Alexa Skill Builder has created a custom skill about basketball including a HowToPlayBasketball intent. When looking at the Intent History page in the developer console, the Builder sees that a number of users are asking the skill how to play baseball. The Builder wants to add a relevant response directing the user back to the topic of basketball.

How should the Builder implement this?

A. Add AMAZON.FallbackIntent and respond with a message about basketball in the handler

B. Create a custom intent related to baseball, and when matched, provide a relevant response

C. Add more sample utterances related to baseball in the HowToPlayBasketball intent

D. Create a new custom baseball slot and add a slot-filling utterance to the HowToPlayBasketball intent.

ANSWER: A

Explanation:

Add `AMAZON.FallbackIntent` and respond with a message about basketball in the handler. This built-in intent is designed to handle customer requests that the skill cannot match to any of its defined intents, including out-of-domain requests such as asking a basketball-focused skill how to play baseball. The skill's fallback-intent handler can provide a concise, helpful redirect, such as explaining that the skill covers basketball and offering examples of supported basketball questions. This gives customers a graceful response instead of an unhandled-request error and preserves the skill's intended scope. The fallback intent should be included in the interaction model and handled in the skill's request-routing code. Amazon recommends using it to recover from unexpected utterances and guide customers toward supported

functionality. See the AMAZON.FallbackIntent documentation and Amazon's guidance on responding gracefully to unexpected requests.