1) An Alexa Skill Builder is writing a skill that needs to play short MP3 files using the SSML audio tag.

Which AWS service provides the MOST highly available and scalable solution to host the MP3 files?

    A) Store the files in an Amazon S3 bucket.
    B) Store the files on an Amazon EC2 instance.
    C) Store the files in an Amazon DynamoDB table.
    D) Store the files in Amazon ElastiCache.

2) When configuring account linking for an Amazon Alexa skill using Login with Amazon, which of the following example scopes should an Alexa Skill Builder use to avoid exposing Personally Identifiable Information (PII)?

    A) profile:email
    B) profile:user_id
    C) profile
    D) profile:name

3) An Alexa Skill Builder is debugging an AWS Lambda function using console.log() statements to inspect values in requests.

How will the log statements be made available to the Builder?

    A) The log files will be saved in Amazon CloudWatch Logs.
    B) A card will be created in the Amazon Alexa app containing the output.
    C) The messages will be available in the console when using ask simulate --debug.
    D) The messages will be spoken as part of the Amazon Alexa output.

4) A skill to search for local businesses returns the following output in the response:

```
<speak>I've found <emphasis level="strong">Bananas & Bears</emphasis></speak>
```

In testing, the Alexa Skill Builder notices the skill keeps failing and reviews the logs in Amazon CloudWatch and finds the following error:

```
"error": {
    "type": "INVALID_RESPONSE" ,
    "message": "Invalid SSML Output Speech for requestId amzn1.echo-
api.request.afb31745-05ad-400e-8703-3a. Error: Fatal error occurred when processing
SSML content. This usually happens when the SSML is not well formed. Error:
Unexpected character ' ' (code 32) (missing name?)\n at [row,col {unknown-source}]:
[1,101]"
}
```

What is wrong with the <speak> response?

A) \<emphasis\> should be \<say-as\>.
B) The entire string needs to be in quotes.
C) The apostrophe in "I've" should be removed.
D) There is an ampersand (&) in the text.

5) What is required for a skill to support multi-turn conversations?

A) The skill must have slots marked "required" with prompt text.
B) The skill must respond to the user by asking a follow-up question.
C) The skill must define session attributes to store the questions and responses.
D) The skill must return shouldEndSession=false in a response.

6) An Alexa Skill Builder is creating a podcast skill using the ASK SDK and AWS Lambda. The Builder wants the skill to launch differently depending on where the user is in the podcast: midway, completed, or at the beginning.

Which type of attribute should be used to store the resume point across skill invocations?

A) Persistence attribute
B) User attribute
C) Session attribute
D) Device attribute

7) To which of the following Amazon Alexa requests can an outputSpeech object be returned? (Select TWO.)

A) IntentRequest
B) SessionEndedRequest
C) Display.ElementSelected
D) PlaybackController.NextCommandIssued
E) AudioPlayer.PlaybackStarted

8) An Alexa Skill Builder has selected a skill invocation name "cat and dog facts."

Why is this an invalid skill invocation name?

A) It has too many characters.
B) It contains more than three syllables.
C) It contains the connecting word "and."
D) It is grammatically incorrect.

9) An Alexa Skill Builder is developing a skill that enables users to get information about train connections between two cities. The skill's invocation name is "train finder" and the main intent for searching connections has the following sample utterances:

"trains between {cityA} and {cityB}"
"connections between {cityA} and {cityB}"
"how do I go from {cityA} to {cityB}"
"when is the next train from {cityA} to {cityB}"

Which example phrases would pass the certification process? (Select TWO.)

A) "Alexa, how can I go from Manchester to London"
B) "Alexa, ask train finder for connections from Manchester to London"
C) "Alexa, ask train finder when is the next train from London to Manchester"
D) "How do I go from Manchester to London"
E) "Alexa, ask train finder what are the trains between London and Manchester"

10) The following JSON defines a custom slot type named `MEDIA_TYPE` with two values:

```
{
  "types": [
    {
      "name": "MEDIA_TYPE",
      "values": [
        {
          "id": "SONG",
          "name": {
            "value": "song",
            "synonyms": ["tune","single","track"]
          }
        },
        {
          "id": "ALBUMS",
          "name": {
            "value": "album",
            "synonyms": ["record","lp","cd"]
          }
        }
      ]
    }
  ]
}
```

Assuming that when a user says "change the record," an intent is matched, and the `MEDIA_TYPE` slot is populated, identify the canonical value that will be returned.

A) ALBUMS
B) album
C) record
D) song

Answers

1) A – S3 is a highly available and durable object store with unlimited storage capacity. It is ideal for hosting the audio files for your skill.

2) B – profile:user_id uniquely identifies the user but does not expose additional information beyond the User ID, which is anonymous.

3) A – Lambda functions save log files to CloudWatch Logs by default. Logging entries are stored in CloudWatch Logs by using standard logging statements such as console.log() in Node.js.

4) D – Ampersands cannot be parsed by SSML and will result in an error when used in speak output in a skill. Ampersands are considered a control character in XML.

5) D – shouldEndSession is a parameter in the Response object that tells Alexa whether or not to keep the session open. True means that the session should end after Alexa speaks the response, and false that the session should remain active. If not provided, it defaults to true. Session attributes are not required when using a dialog model because the user's responses are sent to the skill in the intent's slot values.

6) A – Data that is needed to be retained between sessions is stored in DynamoDB by using the Attributes Persistence Adapter in the Alexa SDK. The concept of Attributes Persistence is discussed in the High/Low Skill Game. Session attributes would not be retained when the user stops the podcast and then resumes later.

7) A,C – An IntentRequest is sent when the user makes a request that corresponds to one of the intents defined in your schema. A Display.ElementSelected request is sent when the user selects an item on a screen on a device that supports it, like Echo Show. You cannot return a response to a SessionEndedRequest. PlaybackController and AudioPlayer requests can only be responded to with an AudioPlayer directive.

8) C – The invocation name must not contain any of the Alexa skill launch phrases and connecting words.

9) C,D – Example phrases must be 100% supported by the sample utterances. The phrase in answer C is fully supported by the sample utterance: "when is the next train from {cityA} to {cityB}." The example phrase in answer D is covered by the sample utterance: "how do I go from {cityA} to {cityB}."

10) B – Even though the user used the word "record" the canonical slot value contains the more general "album" value because "record" is defined as a synonym. ALBUMS is returned as the values id.