


AWS-Certified-Machine-Learning-Specialty

Lernressourcen, AWS-Certified-Machine-Learning-Specialty PDF



Außerdem sind jetzt einige Teile dieser EchteFrage AWS-Certified-Machine-Learning-Specialty Prüfungsfragen kostenlos erhältlich:
https://drive.google.com/open?id=1nP7lfNDrak5xmaf4ae_dw_M2eGRBSSqP

Weil es nicht leicht ist, die Amazon AWS-Certified-Machine-Learning-Specialty Zertifizierungsprüfung zu bestehen. So stellen geeignete Prüfungsmaterialien eine Garantie für den Erfolg dar. EchteFrage wird Ihnen so schnell wie möglich die Amazon AWS-Certified-Machine-Learning-Specialty Prüfungsmaterialien und Fragen und Antworten bieten, so dass Sie sich gut auf die Amazon AWS-Certified-Machine-Learning-Specialty Zertifizierungsprüfung vorbereiten und die Prüfung 100% bestehen können. Mit EchteFrage können Sie nicht nur einmalig AWS-Certified-Machine-Learning-Specialty Prüfung erfolgreich ablegen, sonder auch viel Zeit und Energie ersparen.

Die Amazon AWS-zertifizierte Machine-Learning-Spezifikation (AWS Certified Machine Learnt-Specialty) -Zertifizierungsprüfung soll die Fähigkeiten und Kenntnisse von Fachleuten testen, die mit maschinellen Lerntechnologien in der Amazon Web Services (AWS) Umgebung arbeiten. Diese Zertifizierung ist ideal für Personen, die ihre Fähigkeiten in Bezug auf das Entwerfen, Implementieren und Aufrechterhalten von Lösungen für maschinelles Lernen auf AWS nachweisen möchten. Die Prüfung bewertet Kandidaten zu einer Reihe von Themen, einschließlich Datenentwicklung, Algorithmen für maschinelles Lernen, AWS -Dienste für maschinelles Lernen sowie Modellbereitstellung und -wartung.

>> AWS-Certified-Machine-Learning-Specialty Lernressourcen <<

AWS-Certified-Machine-Learning-Specialty Prüfungsressourcen: AWS Certified Machine Learning - Specialty & AWS-Certified-Machine-Learning-Specialty Reale Fragen

EchteFrage hat riesige Expertenteam, die Ihnen gültige Schulungsressourcen bieten. Sie haben die Amazon AWS-Certified-Machine-Learning-Specialty (AWS Certified Machine Learning - Specialty) Prüfungen in den letzten Jahren nach ihren Erfahrungen und Kenntnissen untersucht. Und endlich kommen die zielgerichteten Fragen und Antworten auf, die den IT-Kandidaten große Hilfe bieten. Nun können Sie im Internet Demo zur Amazon AWS-Certified-Machine-Learning-Specialty (AWS Certified Machine Learning - Specialty) Zertifizierungsprüfung kostenlos herunterladen. Viele IT-Fachleute haben bewiesen, dass EchteFrage sehr zuverlässig ist. Wenn Sie die zielgerichteten Prüfungsfragen von EchteFrage benutzt haben, können Sie normalerweise die Amazon AWS-Certified-Machine-Learning-Specialty Zertifizierungsprüfung bestehen. Schicken Sie doch die Produkte von EchteFrage in den Warenkorb. Sie werden sehr wahrscheinlich der nächste erfolgreiche IT-Fachmann.

Amazon AWS Certified Machine Learning - Specialty AWS-Certified-Machine-Learning-Specialty Prüfungsfragen mit Lösungen (Q138-Q143):

138. Frage

A company has set up and deployed its machine learning (ML) model into production with an endpoint using Amazon SageMaker hosting services. The ML team has configured automatic scaling for its SageMaker instances to support workload changes. During testing, the team notices that additional instances are being launched before the new instances are ready. This behavior needs to change as soon as possible.

How can the ML team solve this issue?

- A. Decrease the cooldown period for the scale-in activity. Increase the configured maximum capacity of instances.
- B. Set up Amazon API Gateway and AWS Lambda to trigger the SageMaker inference endpoint.
- **C. Increase the cooldown period for the scale-out activity.**
- D. Replace the current endpoint with a multi-model endpoint using SageMaker.

Antwort: C

Begründung:

Explanation

The correct solution for changing the scaling behavior of the SageMaker instances is to increase the cooldown period for the scale-out activity. The cooldown period is the amount of time, in seconds, after a scaling activity completes before another scaling activity can start. By increasing the cooldown period for the scale-out activity, the ML team can ensure that the new instances are ready before launching additional instances. This will prevent over-scaling and reduce costs¹ The other options are incorrect because they either do not solve the issue or require unnecessary steps. For example:

Option A decreases the cooldown period for the scale-in activity and increases the configured maximum capacity of instances. This option does not address the issue of launching additional instances before the new instances are ready. It may also cause under-scaling and performance degradation.

Option B replaces the current endpoint with a multi-model endpoint using SageMaker. A multi-model endpoint is an endpoint that can host multiple models using a single endpoint. It does not affect the scaling behavior of the SageMaker instances. It also requires creating a new endpoint and updating the application code to use it² Option C sets up Amazon API Gateway and AWS Lambda to trigger the SageMaker inference endpoint.

Amazon API Gateway is a service that allows users to create, publish, maintain, monitor, and secure APIs. AWS Lambda is a service that lets users run code without provisioning or managing servers.

These services do not affect the scaling behavior of the SageMaker instances. They also require creating and configuring additional resources and services^{3,4} References:

1: Automatic Scaling - Amazon SageMaker

2: Create a Multi-Model Endpoint - Amazon SageMaker

3: Amazon API Gateway - Amazon Web Services

4: AWS Lambda - Amazon Web Services

139. Frage

A developer at a retail company is creating a daily demand forecasting model. The company stores the historical hourly demand data in an Amazon S3 bucket. However, the historical data does not include demand data for some hours.

The developer wants to verify that an autoregressive integrated moving average (ARIMA) approach will be a suitable model for the use case.

How should the developer verify the suitability of an ARIMA approach?

- A. Use Amazon SageMaker Data Wrangler. Import the data from Amazon S3. Resample data by using the aggregate daily total. Perform a Seasonal Trend decomposition.
- B. Use Amazon SageMaker Autopilot. Create a new experiment that specifies the S3 data location. Choose ARIMA as the machine learning (ML) problem. Check the model performance.
- **C. Use Amazon SageMaker Data Wrangler. Import the data from Amazon S3. Impute hourly missing data. Perform a Seasonal Trend decomposition.**
- D. Use Amazon SageMaker Autopilot. Create a new experiment that specifies the S3 data location. Impute missing hourly values. Choose ARIMA as the machine learning (ML) problem. Check the model performance.

Antwort: C

Begründung:

The best solution to verify the suitability of an ARIMA approach is to use Amazon SageMaker Data Wrangler. Data Wrangler is a feature of SageMaker Studio that provides an end-to-end solution for importing, preparing, transforming, featurizing, and analyzing data. Data Wrangler includes built-in analyses that help generate visualizations and data insights in a few clicks. One of the built-in analyses is the Seasonal-Trend decomposition, which can be used to decompose a time series into its trend, seasonal, and residual components. This analysis can help the developer understand the patterns and characteristics of the time series, such as stationarity, seasonality, and autocorrelation, which are important for choosing an appropriate ARIMA model. Data Wrangler also provides built-in transformations that can help the developer handle missing data, such as imputing with mean, median, mode, or constant values, or dropping rows with missing values. Imputing missing data can help avoid gaps and irregularities in the time series, which can affect the ARIMA model performance. Data Wrangler also allows the developer to export the prepared data and the analysis code to various destinations, such as SageMaker Processing, SageMaker Pipelines, or SageMaker Feature Store, for further processing and modeling.

The other options are not suitable for verifying the suitability of an ARIMA approach. Amazon SageMaker Autopilot is a feature-set that automates key tasks of an automatic machine learning (AutoML) process. It explores the data, selects the algorithms relevant to the problem type, and prepares the data to facilitate model training and tuning. However, Autopilot does not support ARIMA as a machine learning problem type, and it does not provide any visualization or analysis of the time series data. Resampling data by using the aggregate daily total can reduce the granularity and resolution of the time series, which can affect the ARIMA model accuracy and applicability.

References:

- *Analyze and Visualize
- *Transform and Export
- *Amazon SageMaker Autopilot
- *ARIMA Model - Complete Guide to Time Series Forecasting in Python

140. Frage

A Data Scientist is working on an application that performs sentiment analysis. The validation accuracy is poor, and the Data Scientist thinks that the cause may be a rich vocabulary and a low average frequency of words in the dataset.

Which tool should be used to improve the validation accuracy?

- A. Natural Language Toolkit (NLTK) stemming and stop word removal
- B. Amazon SageMaker BlazingText cbowmode
- **C. Scikit-learn term frequency-inverse document frequency (TF-IDF) vectorizer**
- D. Amazon Comprehend syntax analysis and entity detection

Antwort: C

Begründung:

Explanation/Reference: <https://monkeylearn.com/sentiment-analysis/>

141. Frage

A Machine Learning Specialist deployed a model that provides product recommendations on a company's website. Initially, the model was performing very well and resulted in customers buying more products on average. However, within the past few months, the Specialist has noticed that the effect of product recommendations has diminished and customers are starting to return to their original habits of spending less. The Specialist is unsure of what happened, as the model has not changed from its initial deployment over a year ago.

Which method should the Specialist try to improve model performance?

- A. The model should be periodically retrained from scratch using the original data while adding a regularization term to handle product inventory changes
- B. The model needs to be completely re-engineered because it is unable to handle product inventory changes.
- C. The model's hyperparameters should be periodically updated to prevent drift.
- **D. The model should be periodically retrained using the original training data plus new data as product inventory changes.**

Antwort: D

142. Frage

An insurance company developed a new experimental machine learning (ML) model to replace an existing model that is in production. The company must validate the quality of predictions from the new experimental model in a production environment before the company uses the new experimental model to serve general user requests.

Which one model can serve user requests at a time. The company must measure the performance of the new experimental model without affecting the current live traffic Which solution will meet these requirements?

- A. A/B testing
- **B. Shadow deployment**
- C. Blue/green deployment
- D. Canary release

Antwort: B

Begründung:

Explanation

The best solution for this scenario is to use shadow deployment, which is a technique that allows the company to run the new experimental model in parallel with the existing model, without exposing it to the end users. In shadow deployment, the company can route the same user requests to both models, but only return the responses from the existing model to the users. The responses from the new experimental model are logged and analyzed for quality and performance metrics, such as accuracy, latency, and resource consumption¹².

This way, the company can validate the new experimental model in a production environment, without affecting the current live traffic or user experience.

The other solutions are not suitable, because they have the following drawbacks:

A: A/B testing is a technique that involves splitting the user traffic between two or more models, and comparing their outcomes based on predefined metrics. However, this technique exposes the new experimental model to a portion of the end users, which might affect their experience if the model is not reliable or consistent with the existing model³.

B: Canary release is a technique that involves gradually rolling out the new experimental model to a small subset of users, and monitoring its performance and feedback. However, this technique also exposes the new experimental model to some end users, and requires careful selection and segmentation of the user groups⁴.

D: Blue/green deployment is a technique that involves switching the user traffic from the existing model (blue) to the new experimental model (green) at once, after testing and verifying the new model in a separate environment. However, this technique does not allow the company to validate the new experimental model in a production environment, and might cause service disruption or inconsistency if the new model is not compatible or stable⁵.

References:

- 1: Shadow Deployment: A Safe Way to Test in Production | LaunchDarkly Blog
- 2: Shadow Deployment: A Safe Way to Test in Production | LaunchDarkly Blog
- 3: A/B Testing for Machine Learning Models | AWS Machine Learning Blog
- 4: Canary Releases for Machine Learning Models | AWS Machine Learning Blog
- 5: Blue-Green Deployments for Machine Learning Models | AWS Machine Learning Blog

143. Frage

.....

Die Amazon AWS-Certified-Machine-Learning-Specialty Dumps von EchteFrage sind die besten Prüfungsunterlagen. Diese Dumps ist unbedingt die Unterlagen, die Sie für länger gesucht haben. Die sind die Prüfungsunterlagen, die speziell für die Prüfungsteilnehmer geschaffen sind. Es kann Ihnen helfen, in sehr kürzer Zeit Amazon AWS-Certified-Machine-Learning-Specialty Zertifizierungsprüfung vorzubereiten und diese Prüfung sehr einfach zu bestehen. Wenn Sie nicht viel Zeit für die Prüfungsvorbereitung, die Amazon AWS-Certified-Machine-Learning-Specialty Dumps von EchteFrage die beste Wahl für sie sind. Damit können Sie Ihre Lerneffektivität verbessern und viel Zeit sparen.

- [illegible]

P.S. Kostenlose 2026 Amazon AWS-Certified-Machine-Learning-Specialty Prüfungsfragen sind auf Google Drive freigegeben von EchteFrage verfügbar: https://drive.google.com/open?id=1nP7lfnDrak5xmaF4ae_dw_M2eGRBSSqP