

Unit 9 - Week 7 : Let Google keep an eye on things. You have the data, but what are you doing with it

Course outline
How does an NPTEL online course work?
Week 0: Welcome to the Course
Week 1: What's the cloud anyway? Start with a solid platform
Week 2: Use GCP to build your apps
Week 3 : Where do I store this Stuff?
Week 4: There's an API for that! You can't secure the Cloud right?
Week 5 : It helps to Network
Week 6 : It helps to Network (continued)
Week 7 : Let Google keep an eye on things. You have the data, but what are you doing with it
<ul style="list-style-type: none"> ● Introduction - Let Google keep an eye on things ● Introduction to IaC ● Cloud Deployment Manager ● Monitoring and Managing Your Services, Apps, and Infra ● Stackdriver [With Lab] ● Recap Quiz - Let Google keep an eye on things ● Summary - Let Google keep an eye on things ● Introduction - You have the data, but what are you doing with it? ● Intro to Big Data Managed Services in the Cloud ● Leverage Big Data Operations with Cloud Dataproc [With Labs] ● Build ETL Pipelines using Cloud Dataflow [With Labs] ● BigQuery Google's Enterprise Data Warehouse ● Recap Quiz - You have the data, but what are you doing with it? ● Summary - You have the data, but what are you doing with it?
○ Quiz : Assignment 7
○ Google Cloud Computing Foundations Course : Week 7 Feedback Form
Let machines do the work
Download Videos

Assignment 7

The due date for submitting this assignment has passed.
As per our records you have not submitted this assignment.

Due on 2020-04-15, 23:59 IST.

- Cloud Deployment Manager allows you to specify all the resources needed for your application in a declarative format using YAML. 1 point

True

False

No, the answer is incorrect.
Score: 0

Feedback:
The answer is True. Cloud Deployment Manager is an Infrastructure as Code (IaC) tool, designed for the automated construction, modification, and destruction of cloud infrastructures. Cloud Deployment Manager allows you to specify all the resources needed for your application in a declarative format using YAML.

Accepted Answers:
True
- What statement best describes infrastructure as code (IaC)? 1 point

Replacing virtual servers with containers.

Using Java scripts to build your systems.

Using Templates to build virtual machines

Automating systems construction using templates

No, the answer is incorrect.
Score: 0

Feedback:
Cloud Deployment Manager is used to automate the building of infrastructure through the use of declarative templates

Accepted Answers:
Automating systems construction using templates
- Which statement accurately describes Google's relationship with open source infrastructure as code (IaC) tools? 1 point

Open source tools are currently in beta.

Open source IaC tools cannot be used with Google

Open source tools can only be used with a valid license

Google has a team of engineers available to support third-party, open source IaC tools

No, the answer is incorrect.
Score: 0

Feedback:
Google has a team of engineers dedicated to ensuring that GCP support is also available for third-party, open source IaC tools.

Accepted Answers:
Google has a team of engineers available to support third-party, open source IaC tools
- Which Stackdriver service provides visibility of the uptime and overall performance of your app? 1 point

Stackdriver Debugger

Stackdriver Logging

Stackdriver Monitoring

Stackdriver Trace

No, the answer is incorrect.
Score: 0

Feedback:
Stackdriver Monitoring provides insight into how your applications are performing.

Accepted Answers:
Stackdriver Monitoring
- What is the Stackdriver service that will store, search, and analyse log files? 1 point

Stackdriver Debugger

Stackdriver Logging

Stackdriver Monitoring

Stackdriver Trace

No, the answer is incorrect.
Score: 0

Feedback:
Stackdriver is a management suite used to provide monitoring, logging, profiling, error reporting, debugging, and tracing services to GCP. The Stackdriver.Logging service allows you to store, search, and analyse log files

Accepted Answers:
Stackdriver Logging
- Stackdriver is a partner product that is easily installed through the GCP Console 1 point

True

False

No, the answer is incorrect.
Score: 0

Feedback:
The answer is False. Stackdriver is a Google solution

Accepted Answers:
False
- You want to automatically perform a test every 5 minutes to ensure your public web servers are up and running. Which of the following is the best option? 1 point

Open Cloud Shell and periodically ping the servers

Write a script and use cron to schedule when it runs

When a web server goes down have it send a log entry to stackdriver to alert on the outage

Create an uptime check in Stackdriver to periodically test if the web servers are reachable

No, the answer is incorrect.
Score: 0

Feedback:
Uptime checks can be configured to periodically test the connection to servers from several remote locations

Accepted Answers:
Create an uptime check in Stackdriver to periodically test if the web servers are reachable
- What is the purpose of Stackdriver Error Reporting? 1 point

To inject errors into an application for testing

To send notifications when an instance crashes

To report on storage issues with persistent disks

To count, analyze, and aggregate crashes in running cloud services in real time

No, the answer is incorrect.
Score: 0

Feedback:
Error Reporting counts, analyzes, and aggregates the crashes in your running cloud services in real time

Accepted Answers:
To count, analyze, and aggregate crashes in running cloud services in real time
- What is the purpose of Stackdriver Trace? 1 point

Find performance bottlenecks in production.

Inject errors in an application so the error can be traced

Send latency statistics to the Stackdriver Error Reporting dashboard

Collect network latency statistics and publish to a latency dashboard.

No, the answer is incorrect.
Score: 0

Feedback:
With Stackdriver Trace, you can inspect detailed latency information for a single request or view aggregate latency for your entire application.

Accepted Answers:
Find performance bottlenecks in production.
- What is the purpose of Stackdriver Debugger? 1 point

Write code tools to interrogate issues.

Investigate the behavior of code in production

Report on bugs and send notifications to users

Slow down an application and step through the code

No, the answer is incorrect.
Score: 0

Feedback:
Stackdriver Debugger allows you to inspect the state of a running application in real time, without stopping or slowing the application down.

Accepted Answers:
Investigate the behavior of code in production
- Which of the following is a managed solution to run Spark, Pig, Hive, and MapReduce in a batch environment with a managed cluster? 1 point

Cloud Dataflow

Cloud Dataprep

Cloud Dataproc

Cloud Runner

No, the answer is incorrect.
Score: 0

Feedback:
Cloud Dataproc is a managed Hadoop Cluster where customers can run their Spark, Hive, Pig, and MapReduce environments

Accepted Answers:
Cloud Dataproc
- You want to make your Cloud Dataproc cluster disposable while keeping your data safe. Which of the following is the best option? 1 point

Leverage persistent disks.

Have the cluster stream data to filestore

Keep the cluster and continue to use HDFS.

Replace your HDFS filesystem with Cloud Storage

No, the answer is incorrect.
Score: 0

Feedback:
Cloud Storage can be used as a replacement for the HDFS. In this example, the cluster can be destroyed without losing data

Accepted Answers:
Replace your HDFS filesystem with Cloud Storage
- You are in need of a service that can process both streaming and batch data, but you don't want to manage the infrastructure. Which of the following is the best solution? 1 point

BigQuery

Cloud Dataflow

Cloud Dataprep

Cloud Dataproc

No, the answer is incorrect.
Score: 0

Feedback:
Cloud Dataflow is a managed service that can process both batch and streaming data

Accepted Answers:
Cloud Dataflow
- "Track" is one of the pipeline components for Cloud Dataflow 1 point

True

False

No, the answer is incorrect.
Score: 0

Feedback:
The answer is False. Cloud Dataflow uses a Source of Information, a series of PCollections that hold data, and a Sink, where the data will be stored or sent. PCollections are the input and output of every transform operation

Accepted Answers:
False
- What code can Cloud Dataflow be written in? 1 point

C++

Java

PHP

Python

Ruby

No, the answer is incorrect.
Score: 0

Feedback:
Cloud Dataflow code can be written in Go, Java, and Python.

Accepted Answers:
Java
Python
- Google provides quick start templates for Cloud Dataflow. Why would you use them? 1 point

To help get started coding.

To make it easier to use Spark

To build pipelines using Cloud Deployment Manager

To rapidly deploy data pipelines without programming expertise

No, the answer is incorrect.
Score: 0

Feedback:
Google provides quick start templates for Cloud Dataflow to allow you to rapidly deploy a number of useful data pipelines without requiring any Apache Beam programming expertise

Accepted Answers:
To rapidly deploy data pipelines without programming expertise
- Select the service that is a managed data warehouse in GCP 1 point

BigQuery

Cloud Bigtable

Cloud Datastore

Cloud SQL

No, the answer is incorrect.
Score: 0

Feedback:
BigQuery is a fully managed data warehouse and analytics engine. The user will never have to build servers or manage storage to make use of BigQuery.

Accepted Answers:
BigQuery
- Which of the following statements regarding BigQuery is accurate? 1 point

BigQuery is a NoOps managed service

BigQuery is a drop-in replacement for Oracle

Before using BigQuery, you have to build a cluster

BigQuery is a SQL database and requires tight schema

No, the answer is incorrect.
Score: 0

Feedback:
BigQuery is a fully managed service. No operations on your part are required.

Accepted Answers:
BigQuery is a NoOps managed service
- What service leverages the Apache Beam SDK to perform ETL operations? 1 point

Cloud Dataflow

Cloud Dataprep

Cloud Dataproc

Cloud Firestore

No, the answer is incorrect.
Score: 0

Feedback:
Cloud Dataflow leverages the Apache Beam SDK to build pipelines to perform ETL operations on batch and streaming data

Accepted Answers:
Cloud Dataflow
- Data can be loaded into BigQuery using the BigQuery Web UI, the BigQuery API, or the BQ command line. 1 point

True

False

No, the answer is incorrect.
Score: 0

Feedback:
The answer is True. BigQuery is a fully managed data warehouse and analytics engine. The user will never have to build servers or manage storage to make use of BigQuery. Loading data into the data warehouse can be easily completed by using the GCP Console, The BigQuery REST based APIs, or through the use of the BigQuery command line (bq) via the SDK

Accepted Answers:
True