

HPE Digital Learner Google Cloud for Architects (Intermediate) Content Pack

HPE Content Pack number	CP019
Content Pack length	22 Hours
Content Pack category	Category 2
Learn more	View now

Why HPE Education Services?

- IDC MarketScape leader 5 years running for IT education and training*
- Recognized by IDC for leading with global coverage, unmatched technical expertise, and targeted education consulting services*
- Key partnerships with industry leaders OpenStack®, VMware®, Linux®, Microsoft®, ITIL, PMI, CSA, and SUSE
- Complete continuum of training delivery options—self-paced eLearning, custom education consulting, traditional classroom, video on-demand instruction, live virtual instructor-led with hands-on lab, dedicated onsite training
- Simplified purchase option with HPE Training Credits

This self-paced eLearning Content Pack provides the required training to help transition from a traditional IT architect role to a cloud architect role utilizing the Google public cloud environment. This is a comprehensive intermediate training series that includes many core areas required to plan, design and implement modern cloud infrastructures, applications and data within the Google public cloud. This training is intended for IT personnel who need to rapidly uplift their IT/cloud skills as part of an ongoing business and operational strategy to migrate IT to the Google public cloud.

Audience

- IT professionals, including managers, engineers and developers, evaluating or implementing application environments on Google Cloud Platform
- Data professionals responsible for provisioning and optimizing big data solutions and data enthusiasts getting started with Google Cloud Platform

Content Pack objectives

This Content Pack provides the information necessary to transition a traditional IT architect into a cloud architect capable of working within a typical Google public cloud platform. This training will introduce many critical baseline knowledge areas needed to transition to a

Google public cloud domain as part of an ongoing hybrid cloud strategy. Areas of focus include Google cloud introductory fundamentals training as well as specific and detailed technology training around a multitude of other relevant topics such as VM management and configuration, containers, compute and applications engine, networking, storage and big data, analytics, GCP scaling, DataProc, APIs and machine learning, identity, access control and security. This training will help the student rapidly transition to a cloud architect role working within a typical Google public cloud environment and will also assist with the path to Google Certification.

Detailed Content Pack outline

Google Cloud

Google Cloud Platform offers a wide array of powerful IaaS and SaaS cloud solutions. In this course, the various service features and options are explored as well as the necessary steps for getting started with Google Cloud Platform with a Google email account.

- Explain the value features of Google Cloud Platform products and services
- Identify the Google Cloud Platform services components, including Compute Engine, App Engine, Container Engine, Container Registry and Cloud Functions
- Identify the Google Cloud Platform storage and database services components, including Cloud Storage, Cloud SQL, Cloud Bigtable, Cloud Datastore and Persistent Disk
- Identify the Google Cloud Platform big data services components, including BigQuery, Cloud Dataflow, Cloud Dataproc, Cloud Datalab, Cloud Pub/Sub and Genomics
- Identify the Google Cloud Platform networking services components, including cloud virtual network, cloud load balancing, cloud CDN, cloud interconnect and cloud DNS
- Define the concepts and components associated with the Google Cloud infrastructure and services
- Define Google Cloud regions with examples
- Define Google Cloud zones with examples
- Differentiate between Cloud Platform's platform as a service (PaaS), software as a service (SaaS) and infrastructure as a service (IaaS)
- Identify the features of Google Cloud Platform's storage and database services
- Identify the features of Google Cloud Platform's networking services
- Identify the features of Google Cloud Platform's big data services
- Manage platform projects, permissions and resources
- Register for a free trial with Google Cloud Platform using a Gmail account
- Create a project using the Google Cloud Platform Console
- Realize options for identity and access management
- Deploy a LAMP stack using Google Cloud Launcher
- Recognize different ways to interact with the Google Cloud Platform
- Recognize the basic components, terms and features of Google Cloud Platform

Google Cloud Virtual Machine Deployment

In this course, you will discover Windows and Linux virtual machine (VM) deployment and explore how to connect to these VM instances for management purposes.

- Describe VM instance configuration details
- Use the GUI to deploy a Windows VM
- Use RDP to remotely manage a Windows VM
- Use the GUI to deploy a Linux VM
- Use SSH to remotely manage a Linux VM

Google Cloud Virtual Machine Configuration

Discover Windows and Linux virtual machine (VM) deployment and explore how to connect to these VM instances for management purposes.

- Describe VM instance configuration details
- Use the GUI to deploy a Windows VM
- Use RDP to remotely manage a Windows VM
- Use the GUI to deploy a Linux VM
- Use SSH to remotely manage a Linux VM

Container, Compute and App Engine with Networking Services

Google Compute Engine offers virtual machines running in Google's innovative data centers and worldwide fiber network. In this course, you will learn about the fundamental aspects of the Google Cloud Platform Container Engine, Compute Engine and App Engine as well as Google's basic networking services.

- Describe Google's Container Engine and Kubernetes
- Define the concept of a container registry
- Deploy an application by creating a container and using the Google Cloud SDK
- Create an image and push it to the Container Registry
- Use kubectl to deploy an application to a container
- Identify characteristics and purpose of Google Compute Engine Networking Service
- Describe the relationship between networks, instances and firewalls
- Create a network in the Web UI
- Create IP addresses in the Web UI
- Create firewall rules in the Web UI
- Describe networks and route collection
- Describe IP forwarding
- Describe Network Address Translation gateways
- Describe network load balancing and HTTP load balancing
- Identify the purpose and benefits of Google App Engine
- Compare traditional builds with Google App Engine builds
- Survey and define the various App Engine services
- Compare the App Engine Standard and App Engine Flexible environments
- Deploy a Python application to the App Engine environment
- Identify GCP attributes of Compute, Container and App Engine as well as networking services

Google Cloud Network Components

Cloud networking borrows from traditional on-premises network configurations. In this course, you will explore and configure Google Cloud network components.

- List Google Cloud network components
- Explain the role that VPCs play in a cloud deployment
- Use the web GUI to manage Google Cloud services
- Deploy a Google cloud VPC with an automatic mode subnet
- Deploy a Google cloud VPC with a manual mode subnet
- Set aside an unchanging public IPv4 address for a VM instance
- Set aside an unchanging public IPv6 address for a VM instance
- Link two Google Cloud VPCs together
- Use the Google Cloud console to add a new subnet
- Explain how Google Cloud firewall rules are processed
- Create the appropriate firewall rules to allow or deny network traffic

Google Cloud Big Data and Machine Learning

Google Cloud Storage is unified object storage for developers and enterprises, from live data serving to data analytics/ML to data archival. In this course, you will learn about the fundamentals of the Google Cloud Datastore and other storage options along with the basics of Big Data and Machine Learning with Google Cloud Platform.

- Identify the purpose and characteristics of Google Cloud Datastore
- Define various datastore terms, including kind, entity, property, keys and entity groups
- Define development libraries, queries, and indexes
- Deploy an App Engine application backed by Google Datastore
- Identify the features of Google Cloud Storage
- Identify the features of Google Cloud SQL
- Identify the features of Google Cloud Bigtable
- Create a Google Cloud Storage bucket and use it to store images
- View objects using the Cloud Storage Browser
- Identify considerations for deployment of Cloud Storage required Applications
- Identify characteristics and purpose of Google Cloud Big Data and Machine Learning platforms
- Recognize loading a CSV File into a BigQuery Table
- Recognize querying data using the CLI
- Recognize querying data using BigQuery Web UI
- Recognize querying data using BigQuery Shell
- Perform interactive queries using BigQuery
- Review the basic features of datastore, storage options, big data and machine learning

Google Cloud Platform Fundamentals

Google Cloud Platform offers several solutions to streamline any enterprise while keeping costs low. This course covers these benefits, including how to navigate GCP and choose between the various data processing products it provides.

- Describe GCP and its use with big data
- Recognize the various services offered by GCP
- Identify the benefits of GCP for data engineers with use case scenarios
- Compare GCP with other models for data engineering
- Perform the steps necessary to create a GCP account
- Describe the Web Admin UI features of the Google Cloud Platform console
- Demonstrate the steps in creating a project
- Demonstrate a GCP process of using BigQuery to query a dataset
- Recognize the various GCP data products and services

Google Cloud Platform Storage and Analytics

Google Cloud Platform (GCP) offers several services for analytics and storage. This course introduces the storage, compute and analytics concepts and how they are tied together.

- Identify the purpose and benefits of using the Compute Engine and Cloud Storage
- Recognize concepts in highly scalable VM management with Compute Engine
- Recognize concepts in each storage option and when to use them
- Demonstrate how to create and manage instances
- Define concepts and features in Cloud Shell
- Learn how to activate and use Google Cloud Shell
- Demonstrate how to create, start and delete a VM instance using 'gcloud beta compute' in Google Cloud Shell
- Demonstrate how to create a Cloud Storage Bucket using 'gsutil mb' in Google Cloud Shell
- Describe GCP concepts in data analytics
- Recognize concepts in Cloud SQL
- Identify the process to create a MySQL database and upload data using Cloud SQL
- Identify the process to create a PostgreSQL database and upload data using Cloud SQL
- Recognize GCP storage, compute and analytics concepts and how they are tied together

Deeper through GCP Analytics and Scaling

The big data industry is getting bigger and GCP has several management tools designed for common use cases. This course introduces common concepts and use cases, including analytics tools and operations.

- Recognize the process for importing data into Cloud SQL
- Work with Apache Hadoop and GCP to manage Hadoop clusters
- Work with the cloud service to run Apache Spark clusters
- Use SparkML for machine learning
- Recognize best practices in big data analysis
- Recall concepts related to the use of GCP Cloud Data Lab
- Describe BigQuery and its benefits in GCP
- Identify datasets with machine learning in GCP
- Recognize machine learning concepts with TensorFlow operations
- Use various machine learning operations to enhance data analysis in common use cases

GCP Network Data Processing Models

You can create and manage an assortment of data processes and network models using GCP. This course will go through the various types, including using a GPU and TensorFlow to create and manage GPU and instances.

- Define the types of virtual networks and the benefits of each
- Specify the process for creating a network
- Recall the process for using TensorFlow with GPU
- Describe the various machine learning APIs and their uses
- Describe Dataflow and how it can be used to create data processing streams
- Recognize differences between Pub and Sub message middleware and when to use them
- Define the various pipelines for Dataflow processing
- Demonstrate the process of creating Dataflow pipelines in GCP
- Specify the differences between real-time and batch data processing
- Recognize more concepts in analysis of data processing in GCP

Google Cloud Data Storage

Cloud storage is easily and rapidly provisioned in Google Cloud. In this module, storage types are discussed along with methods of getting data into the cloud.

- Recognize Google Cloud storage options
- Configure a Google Cloud storage bucket
- Populate a storage bucket with data
- Migrate AWS cloud data to Google Cloud
- List the steps involved in transferring large volumes of data to Google Cloud
- Describe VM instance disk options
- Configure VM instance disks
- Describe characteristics of SQL and NoSQL databases
- Use the GUI to deploy Google Cloud MySQL
- Identify data lifecycle phases
- Configure data retention settings

Google Cloud Dataproc

GCP provides fully managed cloud services for running Apache Spark and Hadoop. This course introduces the concepts of cluster management with Dataproc, including machine types and workers.

- Recognize big data concepts and solutions using GCP
- Define Cloud Dataproc and its benefits
- Recall the various ways to access Dataproc
- Describe the various areas of the dashboard and create a project
- Recognize the process for creating a cluster in Dataproc
- Recall the process for deleting a cluster using Dataproc
- Define master and worker nodes in Dataproc
- Describe custom machine types and preemptible worker nodes
- Define the processes for identity and access management with permissions and IAM roles
- Recognize the basic concepts of cluster management in Dataproc

Dataproc Architecture

Dataproc can be used to perform several operations when integrating platforms, including Pig and Hive. This course will dig further into Dataproc architecture while introducing the use of Pig and Hive.

- Describe how to create a cluster with the Dataproc CLI
- Recognize implementations using the Dataproc REST API
- Describe the various Dataproc architecture types in GCP and common use cases
- Define Dataproc machine types and their uses
- Configure a custom machine type
- Describe how and when to execute Dataproc jobs
- Recognize connections between Apache Hadoop HDFS and Cloud Storage
- Describe the use of Pig and Hive
- Configure and execute a job using Pig and Hive with Dataproc
- Recall concepts of Dataproc jobs, including implementation of Pig and Hive

Continued Dataproc Operations

Executing Dataproc implementations with big data can provide a variety of methods. This course will continue the study of Dataproc implementations with Spark and Hadoop using the cloud shell and introduce BigQuery PySpark REPL package.

- Describe the various Spark and Hadoop processes that can be performed with Dataproc
- Recognize the benefits of separating storage and compute services using Cloud Dataproc
- Recall the process of monitoring and logging Dataproc jobs
- Demonstrate the process of using an SSH tunnel to connect to the master and worker nodes in a cluster
- Define the Spark REPL package and how it is used in Linux
- Describe the compute and storage processes, the benefits of their separation and the virtualized distribution of Hadoop
- Define BigQuery and its benefits for large-scale analytics
- Describe the MapReduce programming model
- Demonstrate the process of submitting multiple jobs with Dataproc
- Recognize the various Dataproc and Cloud Shell job operations and implementations

Implementations with BigQuery for Big Data

You can query big data using the BigQuery tool in the Google Cloud Platform. In this course, you will be introduced to the concepts of using BigQuery, including querying data and using the Google API.

- Describe initialization actions for creating Dataproc clusters
- Specify how functions, operators and data types are used with BigQuery
- Define the steps in installing a storage connector
- Recall the steps to load data with BigQuery
- Identify the steps in exporting and updating data with BigQuery
- Describe the various processes in using Google notebooks with Datalab
- Recall the process for using Jupyter Notebooks with Apache Spark in Google Cloud Dataproc
- Identify the storage options with Google Cloud Storage
- Query data using Google Bigtable
- Identify the various Google APIs
- Describe the concepts of BigQuery and big data using Datalab

APIs and Machine Learning

Machine learning and ML APIs are part of the Cloud ML engine. In this course, you will learn about the various ML APIs and how to implement an app that uses the Vision API.

- Define the cloud ML engine and its purpose with GCP
- Describe the machine learning workflow and how the process works in various scenarios
- Define the Vision API and its purpose
- Describe the Natural Language API and its relevance in GCP
- Recall Translation API and Speech API best practices
- Demonstrate the use of various APIs
- Work with the ML REST API using the Translation API
- Use the Cloud Vision API with a Kubernetes Cluster
- Describe the use of machine learning, including the common APIs

Cloud Identity Management

In this course, you will discover how to create and manage Google Cloud users and groups with Identity and Access Management.

- Describe how user identities are used in the Google Cloud
- Describe how on-premises Active Directory links up with a Google Cloud domain
- Configure Google Cloud IAM users
- Describe how MFA enhances security
- Enable multifactor authentication for a single user
- Enable Cloud Identity for centralized user and group management
- Enable multifactor authentication for Cloud Identity users
- Identify how storage ACLs are managed
- Set storage ACLs to meet business needs

Monitoring and Logging

Cloud resource use must be monitored to ensure proper functionality and use. In this course, you will examine monitoring, alerts and notifications through Stackdriver as well as how to view and export logs.

- Describe the role that SLAs play in cloud computing
- Recognize the importance of monitoring cloud resource management and usage
- Enable Google Cloud resource monitoring using Stackdriver
- Create a custom dashboard
- Recognize the importance of log review
- Use the Google Cloud console to view and export logs
- Use the Google Cloud console to export a VM usage report
- Configure Stackdriver alert policies

Google Cloud Troubleshooting

Troubleshooting problems with deployed cloud services should be done with a structured approach. In this course, you will discover a structured troubleshooting methodology for troubleshooting common problems.

- Describe common troubleshooting steps
- Solve Google Cloud VM disk boot problems
- Solve Google Cloud network connectivity issues
- Solve Google Cloud VPN problems
- Solve Google Cloud permissions issues

Cloud Solution Management and Testing

Cloud solutions should follow a structured approach throughout their lifetime. In this course, you will explore SDLC phases, ITIL and various testing techniques.

- Describe the meaning of each SDLC phase
- Recognize how ITIL provides efficient and cost effective IT service delivery
- Recall how continuous integration and deployment provide updates
- Describe how fuzz testing is used to ensure security and quality
- Describe the purpose of regression testing
- Describe how unit testing uses a modular testing approach
- Specify how load and stress tests can identify capacity weaknesses within a solution

CLI Cloud Resource Management

The gcloud and gsutil CLI tools can be used to deploy and manage Google Cloud resources and automate repetitive administrative tasks. In this course, you will explore these Google Cloud tools.

- Describe how gcloud is used to manage cloud resources
- Take a snapshot using gcloud
- Deploy a VM using gcloud
- Use gcloud to delete a virtual machine instance
- Use gcloud to manage DNS
- Use gcloud to manage firewall settings
- Use gsutil to create a cloud storage bucket
- Use gsutil to copy data into a bucket
- Use gsutil to manage storage ACLs

Google Cloud Programmatic Access

Automation is the key to running repetitive management tasks. In this course, you will explore Google APIs, PowerShell cmdlets and the gcloud command line tool.

- Describe how Google APIs and the SDK allow programmatic access
- Install Google Cloud SDK components on a Windows computer
- List examples of how the Google Cloud Shell allows command line access to resource management
- Use Google Cloud Shell to show VM instances
- Describe the PowerShell syntax and use
- Perform common administrative tasks using Google Cloud PowerShell cmdlets

Google Cloud Web Applications and Name Resolution

In this course, you will explore web application, content delivery networks and DNS name resolution. You will also learn about cloud network connections through VPNs and dedicated links.

- Recognize the role that Google Cloud App Engine plays
- Deploy a Google Cloud web application
- Describe the role DNS plays within a TCP/IP network
- Deploy a Google Cloud DNS configuration
- Describe how content delivery networks improve the end user experience
- Use the Google Cloud Console to enable a content delivery configuration
- Recognize how VPNs are used with cloud computing
- Describe how Dedicated Interconnect provides a private network link to the Google Cloud

GCP Engineering and Streaming Architecture

Feature engineering can be an essential tool in applied machine learning when enhancing a dataset. In this course, you will learn about concepts of feature engineering, including areas of streaming architecture and implementations.

- Describe the concepts of feature engineering
- Recall the benefits of quality features with feature engineering and feature selection
- Describe the process of input selection in feature engineering
- Demonstrate feature engineering in use cases
- Recall the concepts of streaming data and real-time stream processing
- Describe Dataflow triggers and late data
- Install Java JDK on Windows 10
- Demonstrate how to install Apache Maven on Windows 10
- Install Google Cloud SDK and initialize SDK Shell on Windows 10
- Demonstrate the process of streaming pipelines using Dataflow SDK 2.x and Java in Cloud SDK Shell
- Demonstrate the process of streaming pipelines using Dataflow SDK 2.x and Python in Google Cloud Shell
- Describe feature engineering concepts and streaming data architecture

GCP Big Data and Security

Complex operations require more managed and secure systems. This course explores the use of Cloud Spanner and Bigtable to create more complex data service configurations and manage secure data infrastructure.

- Describe Cloud Spanner and its purpose in GCP
- Describe Cloud Spanner replication and instances
- Describe Cloud Spanner schema, datatypes, and best practices
- Describe GCP Bigtable
- Specify the steps to design a Bigtable schema
- Demonstrate the process of creating a Bigtable instance on GCP
- Describe the cloud platform security model
- Describe the use of Cloud Identity and Access Management
- Describe the security layers on GCP
- Define the security standards Google works to comply with
- Describe the use of big data and how to keep data secure on GCP

Learn more at

www.hpe.com/ww/digitallearner

www.hpe.com/ww/digitallearner-contentpack

Follow us:



© Copyright 2019 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

CP019 A.00, January 2019