Networking in Google Cloud Platform

Duration: 2 Days Course Code: GO5976

Overview:

This 2-day instructor-led Google Cloud Platform course gives participants a broad study of networking options on Google Cloud Platform. Through presentations, demonstrations, and hands-on labs, participants explore and deploy GCP networking technologies, such as Google Virtual Private Cloud (VPC) networks, subnets, firewalls; interconnection among networks; load balancing; Cloud DNS; Cloud CDN; Cloud NAT. The course will also cover common network design patterns and automated deployment using Deployment Manager or Terraform.

Target Audience:

This course is intended for the following participants: Network engineers and Admins who are either using Google Cloud Platform or are planning to do soIndividuals who want to be exposed to software-defined networking solutions in the cloud

Objectives:

- This course teaches participants the following skills:
- Configure Google VPC networks, subnets, and routers
- Control administrative access to VPC objects
- Control network access to endpoints in VPCs
- Interconnect networks among GCP projects
- Interconnect networks among GCP VPC networks and on-premises or other-cloud networks
- Choose among GCP load balancer and proxy options and configure them
- Use Cloud CDN to reduce latency and save money
- Optimize network spend using Network Tiers
- Configure Cloud NAT or Private Google Access to provide instances without public IP addresses access to other services
- Deploy networks declaratively using Cloud Deployment Manager or Terraform
- Design networks to meet common customer requirements
- Configure monitoring and logging to troubleshoot networks problems

Prerequisites:

To get the most out of this course, participants should:

- Completed Google Cloud Platform Fundamentals: Core Infrastructure or have equivalent experience
- Prior understanding of the 7 layer OSI model
- Prior understanding of IPv4 addressing
- Prior experience with managing IPv4 routes

Content:

Module 1: Google Cloud VPC Networking Fundamentals

Topics Covered:

- Recall that networks belong to projects
- Explain the differences among default, auto, and custom networks
- Create networks and subnets
- Explain how IPv4 addresses are assigned to Compute Engine instances
- Publish domain names using Google Cloud DNS
- Create Compute Engine instances with IP aliases
- Create Compute Engine instances with multiple virtual network interfaces
- Outline how IAM policies affect VPC networks
- Control access to network resources using service accounts
- Control access to Compute Engine instances with tag-based firewall rules
- Outline the overall workflow for configuring shared VPC
- Differentiate between the IAM roles that allow network resources to be managed
- Configure peering between unrelated VPC networks
- Recall when to use Shared VPC and when to use VPC Network Peering
- Recall the various load balancing services
- Configure Layer 7 HTTP(S) load balancing
- Whitelist and blacklist IP traffic with Cloud Armor
- Cache content with Cloud CDN
- Explain Layer 4 TCP or SSL proxy load balancing
- Explain regional network load balancing
- Configure internal load balancing
- Recall the choices for enabling IPv6 Internet connectivity for GCP load balancers
- Determine which GCP load balancer to use when
- Recall the GCP interconnect and peering services available to connect your infrastructure to GCP
- Explain Dedicated Interconnect and Partner Interconnect
- Describe the workflow for configuring a Dedicated Interconnect
- Build a connection over a VPN with Cloud Router
- Determine which GCP interconnect service to use when
- Explain Direct Peering and Partner Peering
- Determine which GCP peering service to use when
- Recognize how networking features are charged for
- Use Network Service Tiers to optimize spend

Module 4: Load Balancing

Topics Covered:

- Recall that networks belong to projects
- Explain the differences among default, auto, and custom networks
- Create networks and subnets
- Explain how IPv4 addresses are assigned to Compute Engine instances
- Publish domain names using Google Cloud DNS
- Create Compute Engine instances with IP aliases
- Create Compute Engine instances with multiple virtual network interfaces
- Outline how IAM policies affect VPC networks
- Control access to network resources using service accounts
- Control access to Compute Engine instances with tag-based firewall rules
- Outline the overall workflow for configuring shared VPC
- Differentiate between the IAM roles that allow network resources to be managed
- Configure peering between unrelated VPC networks
- Recall when to use Shared VPC and when to use VPC Network Peering
- Recall the various load balancing services
- Configure Layer 7 HTTP(S) load balancing
- Whitelist and blacklist IP traffic with Cloud Armor
- Cache content with Cloud CDN
- Explain Layer 4 TCP or SSL proxy load balancing
- Explain regional network load balancing
- Configure internal load balancing
- Recall the choices for enabling IPv6 Internet connectivity for GCP load balancers
- Determine which GCP load balancer to use when
- Recall the GCP interconnect and peering services available to connect your infrastructure to GCP
- Explain Dedicated Interconnect and Partner Interconnect
- Describe the workflow for configuring a Dedicated Interconnect
- Build a connection over a VPN with Cloud Router
- Determine which GCP interconnect service to use when
- Explain Direct Peering and Partner Peering
- Determine which GCP peering service to use when
- Recognize how networking features are

Module 7: Network Design and Deployment

Topics Covered:

- Recall that networks belong to projects
- Explain the differences among default, auto, and custom networks
- Create networks and subnets
- Explain how IPv4 addresses are assigned to Compute Engine instances
- Publish domain names using Google Cloud DNS
- Create Compute Engine instances with IP aliases
- Create Compute Engine instances with multiple virtual network interfaces
- Outline how IAM policies affect VPC networks
- Control access to network resources using service accounts
- Control access to Compute Engine instances with tag-based firewall rules
- Outline the overall workflow for configuring shared VPC
- Differentiate between the IAM roles that allow network resources to be managed
- Configure peering between unrelated VPC networks
- Recall when to use Shared VPC and when to use VPC Network Peering
- Recall the various load balancing services
- Configure Layer 7 HTTP(S) load balancing
- Whitelist and blacklist IP traffic with Cloud Armor
- Cache content with Cloud CDN
- Explain Layer 4 TCP or SSL proxy load balancing
- Explain regional network load balancing
- Configure internal load balancing
- Recall the choices for enabling IPv6 Internet connectivity for GCP load balancers
- Determine which GCP load balancer to use when
- Recall the GCP interconnect and peering services available to connect your infrastructure to GCP
- Explain Dedicated Interconnect and Partner Interconnect
- Describe the workflow for configuring a Dedicated Interconnect
- Build a connection over a VPN with Cloud Router
- Determine which GCP interconnect service to use when
- Explain Direct Peering and Partner Peering
- Determine which GCP peering service to use when
- Recognize how networking features are charged for
- Use Network Service Tiers to optimize spend

- Determine which Network Service Tier to use when
- Recall that labels can be used to understand networking spend
- Explain common network design patterns
- Configure Private Google Access to allow access to certain Google Cloud services from VM instances with only internal IP addresses
- Configure Cloud NAT to provide your instances without public IP addresses access to the internet
- Automate the deployment of networks using Deployment Manager or Terraform
- Launch networking solutions using Cloud Marketplace
- Configure uptime checks, alerting policies and charts for your network services
- Use VPC Flow Logs to log and analyze network traffic behavior

Module 2: Controlling Access to VPC Networks

Topics Covered:

- Recall that networks belong to projects
- Explain the differences among default, auto, and custom networks
- Create networks and subnets
- Explain how IPv4 addresses are assigned to Compute Engine instances
- Publish domain names using Google Cloud DNS
- Create Compute Engine instances with IP aliases
- Create Compute Engine instances with multiple virtual network interfaces
- Outline how IAM policies affect VPC networks
- Control access to network resources using service accounts
- Control access to Compute Engine instances with tag-based firewall rules
- Outline the overall workflow for configuring shared VPC
- Differentiate between the IAM roles that allow network resources to be managed
- Configure peering between unrelated VPC networks
- Recall when to use Shared VPC and when to use VPC Network Peering
- Recall the various load balancing services
- Configure Layer 7 HTTP(S) load balancing
- Whitelist and blacklist IP traffic with Cloud Armor
- Cache content with Cloud CDN
- Explain Layer 4 TCP or SSL proxy load balancing
- Explain regional network load balancing
- Configure internal load balancing
- Recall the choices for enabling IPv6 Internet connectivity for GCP load balancers
- Determine which GCP load balancer to use when charged for
- Use Network Service Tiers to optimize spend
- Determine which Network Service Tier to use when
- Recall that labels can be used to understand networking spend
- Explain common network design patterns
- Configure Private Google Access to allow access to certain Google Cloud services from VM instances with only internal IP addresses
- Configure Cloud NAT to provide your instances without public IP addresses access to the internet
- Automate the deployment of networks using Deployment Manager or Terraform
- Launch networking solutions using Cloud Marketplace
- Configure uptime checks, alerting policies and charts for your network services
- Use VPC Flow Logs to log and analyze network traffic behavior

Module 5: Hybrid Connectivity

Topics Covered:

- Recall that networks belong to projects
- Explain the differences among default, auto, and custom networks
- Create networks and subnets
- Explain how IPv4 addresses are assigned to Compute Engine instances
- Publish domain names using Google Cloud DNS
- Create Compute Engine instances with IP aliases
- Create Compute Engine instances with multiple virtual network interfaces
- Outline how IAM policies affect VPC networks
- Control access to network resources using service accounts
- Control access to Compute Engine instances with tag-based firewall rules
- Outline the overall workflow for configuring shared VPC
- Differentiate between the IAM roles that allow network resources to be managed
- Configure peering between unrelated VPC networks
- Recall when to use Shared VPC and when to use VPC Network Peering
- Recall the various load balancing services
- Configure Layer 7 HTTP(S) load balancing
- Whitelist and blacklist IP traffic with Cloud Armor
- Cache content with Cloud CDN
- Explain Layer 4 TCP or SSL proxy load balancing
- Explain regional network load balancing
- Determine which Network Service Tier to use when
- Recall that labels can be used to understand networking spend
- Explain common network design patterns
- Configure Private Google Access to allow access to certain Google Cloud services from VM instances with only internal IP addresses
- Configure Cloud NAT to provide your instances without public IP addresses access to the internet
- Automate the deployment of networks using Deployment Manager or Terraform
- Launch networking solutions using Cloud Marketplace
- Configure uptime checks, alerting policies and charts for your network services
- Use VPC Flow Logs to log and analyze network traffic behavior

Module 8: Network Monitoring and Troubleshooting

Topics Covered:

- Recall that networks belong to projects
- Explain the differences among default, auto, and custom networks
- Create networks and subnets
- Explain how IPv4 addresses are assigned to Compute Engine instances
- Publish domain names using Google Cloud DNS
- Create Compute Engine instances with IP aliases
- Create Compute Engine instances with multiple virtual network interfaces
- Outline how IAM policies affect VPC networks
- Control access to network resources using service accounts
- Control access to Compute Engine instances with tag-based firewall rules
- Outline the overall workflow for configuring shared VPC
- Differentiate between the IAM roles that allow network resources to be managed
- Configure peering between unrelated VPC networks
- Recall when to use Shared VPC and when to use VPC Network Peering
- Recall the various load balancing services
- Configure Layer 7 HTTP(S) load balancing
- Whitelist and blacklist IP traffic with Cloud Armor
- Cache content with Cloud CDN
- Explain Layer 4 TCP or SSL proxy load balancing
- Explain regional network load balancing
- Configure internal load balancing
- Recall the choices for enabling IPv6 Internet connectivity for GCP load balancers

- Recall the GCP interconnect and peering services available to connect your infrastructure to GCP
- Explain Dedicated Interconnect and Partner Interconnect
- Describe the workflow for configuring a Dedicated Interconnect
- Build a connection over a VPN with Cloud Router
- Determine which GCP interconnect service to use when
- Explain Direct Peering and Partner Peering
- Determine which GCP peering service to use when
- Recognize how networking features are charged for
- Use Network Service Tiers to optimize spend
- Determine which Network Service Tier to use when
- Recall that labels can be used to understand networking spend
- Explain common network design patterns
- Configure Private Google Access to allow access to certain Google Cloud services from VM instances with only internal IP addresses
- Configure Cloud NAT to provide your instances without public IP addresses access to the internet
- Automate the deployment of networks using Deployment Manager or Terraform
- Launch networking solutions using Cloud Marketplace
- Configure uptime checks, alerting policies and charts for your network services
- Use VPC Flow Logs to log and analyze network traffic behavior

Module 3: Sharing Networks across Projects

Topics Covered:

- Recall that networks belong to projects
- Explain the differences among default, auto, and custom networks
- Create networks and subnets
- Explain how IPv4 addresses are assigned to Compute Engine instances
- Publish domain names using Google Cloud DNS
- Create Compute Engine instances with IP aliases
- Create Compute Engine instances with multiple virtual network interfaces
- Outline how IAM policies affect VPC networks
- Control access to network resources using service accounts
- Control access to Compute Engine instances with tag-based firewall rules
- Outline the overall workflow for configuring shared VPC
- Differentiate between the IAM roles that
- Configure internal load balancing
- Recall the choices for enabling IPv6 Internet connectivity for GCP load balancers
- Determine which GCP load balancer to use when
- Recall the GCP interconnect and peering services available to connect your infrastructure to GCP
- Explain Dedicated Interconnect and Partner Interconnect
- Describe the workflow for configuring a Dedicated Interconnect
- Build a connection over a VPN with Cloud Router
- Determine which GCP interconnect service to use when
- Explain Direct Peering and Partner Peering
- Determine which GCP peering service to use when
- Recognize how networking features are charged for
- Use Network Service Tiers to optimize spend
- Determine which Network Service Tier to use when
- Recall that labels can be used to understand networking spend
- Explain common network design patterns
- Configure Private Google Access to allow access to certain Google Cloud services from VM instances with only internal IP addresses
- Configure Cloud NAT to provide your instances without public IP addresses access to the internet
- Automate the deployment of networks using Deployment Manager or Terraform
- Launch networking solutions using Cloud Marketplace
- Configure uptime checks, alerting policies and charts for your network services
- Use VPC Flow Logs to log and analyze network traffic behavior

Module 6: Networking Pricing and Billing

Topics Covered:

- Recall that networks belong to projects
- Explain the differences among default, auto, and custom networks
- Create networks and subnets
- Explain how IPv4 addresses are assigned to Compute Engine instances
- Publish domain names using Google Cloud DNS
- Create Compute Engine instances with IP aliases
- Create Compute Engine instances with multiple virtual network interfaces
- Outline how IAM policies affect VPC networks
- Determine which GCP load balancer to use when
- Recall the GCP interconnect and peering services available to connect your infrastructure to GCP
- Explain Dedicated Interconnect and Partner Interconnect
- Describe the workflow for configuring a Dedicated Interconnect
- Build a connection over a VPN with Cloud Router
- Determine which GCP interconnect service to use when
- Explain Direct Peering and Partner Peering
- Determine which GCP peering service to use when
- Recognize how networking features are charged for
- Use Network Service Tiers to optimize spend
- Determine which Network Service Tier to use when
- Recall that labels can be used to understand networking spend
- Explain common network design patterns
- Configure Private Google Access to allow access to certain Google Cloud services from VM instances with only internal IP addresses
- Configure Cloud NAT to provide your instances without public IP addresses access to the internet
- Automate the deployment of networks using Deployment Manager or Terraform
- Launch networking solutions using Cloud Marketplace
- Configure uptime checks, alerting policies and charts for your network services
- Use VPC Flow Logs to log and analyze network traffic behavior

allow network resources to be managed
- Configure peering between unrelated VPC networks
- Recall when to use Shared VPC and when to use VPC Network Peering
- Recall the various load balancing services
- Configure Layer 7 HTTP(S) load balancing
- Whitelist and blacklist IP traffic with Cloud Armor
- Cache content with Cloud CDN
- Explain Layer 4 TCP or SSL proxy load balancing
- Explain regional network load balancing
- Configure internal load balancing
- Recall the choices for enabling IPv6 Internet connectivity for GCP load balancers
- Determine which GCP load balancer to use when
- Recall the GCP interconnect and peering services available to connect your infrastructure to GCP
- Explain Dedicated Interconnect and Partner Interconnect
- Describe the workflow for configuring a Dedicated Interconnect
- Build a connection over a VPN with Cloud Router
- Determine which GCP interconnect service to use when
- Explain Direct Peering and Partner Peering
- Determine which GCP peering service to use when
- Recognize how networking features are charged for
- Use Network Service Tiers to optimize spend
- Determine which Network Service Tier to use when
- Recall that labels can be used to understand networking spend
- Explain common network design patterns
- Configure Private Google Access to allow access to certain Google Cloud services from VM instances with only internal IP addresses
- Configure Cloud NAT to provide your instances without public IP addresses access to the internet
- Automate the deployment of networks using Deployment Manager or Terraform
- Launch networking solutions using Cloud Marketplace
- Configure uptime checks, alerting policies and charts for your network services
- Use VPC Flow Logs to log and analyze network traffic behavior
- Control access to network resources using service accounts
- Control access to Compute Engine instances with tag-based firewall rules
- Outline the overall workflow for configuring shared VPC
- Differentiate between the IAM roles that allow network resources to be managed
- Configure peering between unrelated VPC networks
- Recall when to use Shared VPC and when to use VPC Network Peering
- Recall the various load balancing services
- Configure Layer 7 HTTP(S) load balancing
- Whitelist and blacklist IP traffic with Cloud Armor
- Cache content with Cloud CDN
- Explain Layer 4 TCP or SSL proxy load balancing
- Explain regional network load balancing
- Configure internal load balancing
- Recall the choices for enabling IPv6 Internet connectivity for GCP load balancers
- Determine which GCP load balancer to use when
- Recall the GCP interconnect and peering services available to connect your infrastructure to GCP
- Explain Dedicated Interconnect and Partner Interconnect
- Describe the workflow for configuring a Dedicated Interconnect
- Build a connection over a VPN with Cloud Router
- Determine which GCP interconnect service to use when
- Explain Direct Peering and Partner Peering
- Determine which GCP peering service to use when
- Recognize how networking features are charged for
- Use Network Service Tiers to optimize spend
- Determine which Network Service Tier to use when
- Recall that labels can be used to understand networking spend
- Explain common network design patterns
- Configure Private Google Access to allow access to certain Google Cloud services from VM instances with only internal IP addresses
- Configure Cloud NAT to provide your instances without public IP addresses access to the internet
- Automate the deployment of networks using Deployment Manager or Terraform
- Launch networking solutions using Cloud Marketplace
- Configure uptime checks, alerting policies and charts for your network services

- Use VPC Flow Logs to log and analyze network traffic behavior

Further Information:

For More information, or to book your course, please call us on 0800/84.009

info@globalknowledge.be

www.globalknowledge.com/en-be/