Google Cloud

Professional Cloud Network Engineer

Certification exam guide

A Professional Cloud Network Engineer is responsible for the design, implementation, and management of Google Cloud network infrastructure. This includes designing network architectures for high availability, scalability, resiliency, and security. This individual is skilled in configuring and managing Virtual Private Cloud (VPC), routing, load balancing, Cloud NAT, and Cloud DNS. They are proficient in setting up hybrid and multicloud connectivity through Cloud Interconnect and Cloud VPN, and their expertise extends to diagnosing, monitoring, and troubleshooting network operations by using Google Cloud Observability and Network Intelligence Center. Additionally, they are experienced in configuring and architecting network security for Google Cloud, including but not limited to securing VPCs and firewalls (Cloud Next Generation Firewall and VPC firewalls), Secure Web Proxy, VPC Service Controls, and Google Cloud Armor.

Section 1: Designing and planning a Google Cloud VPC network (~21% of the exam)

1.1 Designing an overall network architecture. Considerations include:

- Differentiating between network tiers (e.g., Premium and Standard).
- Designing for high availability, failover, disaster recovery, and scale.
- Designing the DNS topology (e.g., on-premises and Cloud DNS).
- Choosing an appropriate load balancer for network implementation.
- Planning for Google Kubernetes Engine (GKE) networking (e.g., secondary ranges, scale potential based on IP address space, and access to GKE control plane).
- Identifying the most appropriate Identity and Access Management (IAM) roles suited to specific network architecture designs (e.g. load balancer provisioning and Shared VPC subnet permissions).
- Planning for connectivity to managed services (e.g., private services access, Private Service Connect [PSC], and Serverless VPC Access).
- Planning for quotas and limits.

1.2 Designing VPC networks. Considerations include:

- Choosing the VPC type and quantity (e.g., standalone or Shared VPC and the number of VPC environments).
- Determining how the networks interconnect based on requirements (e.g., VPC Network Peering, network connectivity [mesh and star topology] with Network Connectivity Center, and PSC).
- Planning the IP address management (IPAM) strategy (e.g., subnets, IPv6, bring your own IP, privately used public IP [PUPI], Private NAT, non-RFC 1918 addresses, managed services, and IPAM automation techniques).
- Planning a global or regional network environment (or variations of these).
- Determining the correct maximum transmission unit (MTU) sizing for VPC for workloads.
- Planning third-party device insertion (e.g., network virtual appliance) with custom routes (static or policy-based) and load balancing.

1.3 Designing a resilient and performant hybrid and multi-cloud network. Considerations include:

- Designing for hybrid (e.g., on-premises and cloud, branch office) connectivity, including bandwidth and security constraints (e.g., Dedicated Interconnect, Partner Interconnect, Cloud VPN, and SD-WAN appliances).
- Designing for multicloud connectivity (e.g., Cloud VPN and Cross-Cloud Interconnect).
- Choosing when to use Direct Peering or Verified Peering Provider.
- Designing high-availability and disaster recovery connectivity strategies for multiple regions (e.g., regional or global dynamic routing mode).
- Accessing multiple VPCs from on-premises locations (e.g., Shared VPC, multi-VPC peering, and Network Connectivity Center topologies).
- Accessing Google services like Vertex AI and application programming interfaces (APIs) privately from on-premises locations.
- Accessing managed services through PSC and VPC Network Peering connections (e.g., private services access).
- Designing the IP address space across on-premises locations and cloud environments (e.g., internal ranges, planning to avoid overlaps, and Private NAT).
- Architecting hybrid DNS topology: Define forwarding paths, inbound policies, cross-project binding, and DNS peering strategy.
- Determining the correct MTU sizing for hybrid connections (Cloud Interconnect and HA VPN) for workloads.
- Understanding interconnect encryption options, such as MACsec and HA VPN, over Cloud Interconnect.

1.4 Designing for Google Kubernetes Engine (GKE). Considerations include:

- Choosing between public or private cluster nodes and node pools.
- Choosing between public or private control plane endpoints.
- Planning subnets: Primary and secondary ranges.
- Planning for GKE IP addresses using (RFC 1918, non-RFC 1918, Google-managed services range, PSC, shared IP ranges, and PUPI).
- Planning for IPv6.
- Designing load balancing for GKE networking.
- Adding and managing node pool configuration.

Section 2: Implementing a VPC network (~20% of the exam)

2.1 Configuring VPCs. Considerations include::

- Creating Google Cloud VPC resources (e.g., networks, subnets, firewall rules or policies, private services access subnet, and private pools).
- Configuring VPC Network Peering.
- Creating a Shared VPC network and sharing subnets with other projects.
- Assigning the correct IAM permissions to use Shared VPC subnets from service projects.
- Configuring access to Google APIs and Google-managed services (e.g., Private Google Access and public interfaces).
- Expanding VPC subnet ranges after creation.
- Configuring restricted Google Cloud services with VPC Service Controls perimeters.

2.2 Configuring VPC routing. Considerations include:

- Setting up static and dynamic routing (e.g., Cloud Router).
- Configuring global or regional dynamic routing.
- Implementing routing using network tags and priority.
- Implementing route priorities with global dynamic routing, including policy-based routing and dynamic routing.
- Implementing an internal load balancer as a next hop.
- Configuring custom route import/export over VPC Network Peering and Network Connectivity Center.
- Configuring policy-based routing.

2.3 Configuring Network Connectivity Center. Considerations include:

- Differentiating between spoke types (VPC spoke, hybrid spoke, and producer spoke).

- Managing VPC topology (e.g., star topology, hub and spokes, and mesh topology).
- Configuring Private NAT and PSC propagation.
- Configuring IP/CIDR range filters for Network Connectivity Center spokes.
- Monitoring and troubleshooting Network Connectivity Center.

2.4 Configuring and maintaining GKE clusters. Considerations include:

- Creating VPC-native clusters using alias IPs.
- Setting up clusters with Shared VPC.
- Configuring private clusters and private control plane endpoints.
- Adding authorized networks for cluster control plane endpoints.
- Using DNS-based endpoint for control plane access.
- Enabling GKE Dataplane V2.
- Configuring source NAT (SNAT) and IP Masquerade policies.
- Creating GKE network policies.
- Configuring Pod ranges and service ranges.
- Deploying additional Pod ranges for GKE clusters.
- Configuring DNS (local DNS cache, Cloud DNS, and kube-dns).

Section 3: Configuring managed network services (~16% of the exam)

3.1 Configuring load balancing. Considerations include:

- Determining the load balancing solution for your network (internal/external, regional/global, application/proxy/passthrough, etc.).
- Configuring backend services, including autoscaling (e.g., network endpoint groups [NEGs] and managed instance groups).
- Configuring various load balancers and backend settings, such as the balancing method, session affinity, serving capacity, URL maps, health checks, and global access.
- Understanding load balancers in GKE (e.g., GKE Gateway controller, GKE Ingress controller, and NEGs).
- Setting up traffic management on Application Load Balancer (e.g., traffic splitting, traffic mirroring, and URL rewrites).

3.2 Configuring Cloud CDN. Considerations include:

- Setting up Cloud CDN for supported origins (e.g., managed instance groups, Cloud Storage buckets, and Cloud Run).

- Setting up Cloud CDN for external backends (internet NEGs) and third-party object storage.
- Invalidating cached content.

3.3 Configuring Cloud DNS. Considerations include:

- Managing Cloud DNS zones and records.
- Migrating to Cloud DNS.
- Configuring Cloud DNS routing policies, such as geolocation and failover policies.
- Enabling DNS Security Extensions (DNSSEC).
- Setting up self-hosted DNS integration with Cloud DNS, including configuring DNS forwarding and DNS server policies.
- Understanding DNS private and public zones and setting up split-horizon DNS.
- Setting up DNS cross-project binding and DNS peering.
- Configuring Cloud DNS and external-DNS operator for GKE.

Section 4: Configuring and implementing hybrid and multicloud network interconnectivity (~16% of the exam)

4.1 Configuring Cloud Interconnect. Considerations include:

- Creating Dedicated Interconnect connections and configuring VLAN attachments.
- Creating Partner Interconnect connections, configuring VLAN attachments, and differentiating between layer 2 and layer 3 type interconnects.
- Creating Cross-Cloud Interconnect connections and configuring VLAN attachments.
- Configuring HA VPN over Cloud Interconnect.
- Implementing 99.9% and 99.99% service-level agreements (SLAs) for interconnect topologies.

4.2 Configuring a site-to-site IPSec VPN. Considerations include:

- Configuring HA VPN toward on-premise VPN gateways.
- Configuring HA VPN toward other Google Cloud VPCs.
- Configuring Classic VPN (e.g., route-based and policy-based).

4.3 Configuring Cloud Router. Considerations include:

- Implementing Border Gateway Protocol (BGP) attributes (e.g., ASN, route priority/MED, link-local addresses, and authentication).

- Configuring Bidirectional Forwarding Detection (BFD).
- Creating custom-advertised routes and custom-learned routes.
- Selecting between legacy and standard best path selection at the VPC.

4.4 Configuring Network Connectivity Center. Considerations include:

- Creating hybrid spokes (e.g., VPN and VLAN attachment).
- Establishing site-to-site data transfer.
- Creating router appliances (RAs).
- Solving common transitivity networking issues.

Section 5: Managing, monitoring, and troubleshooting network operations (~14% of the exam)

5.1 Logging and monitoring with Google Cloud Observability. Considerations include:

- Enabling and reviewing Cloud Logging for networking components (e.g., Cloud VPN, Cloud Router, VPC Service Controls, Cloud Next Generation Firewall [NGFW], Firewall Insights, VPC Flow Logs, Cloud DNS, Cloud NAT, and Network Connectivity Center).
- Monitoring networking metrics (e.g., Cloud VPN, Cloud Interconnect and VLAN attachments, Cloud Router, load balancers, Google Cloud Armor, and Cloud NAT).

5.2 Maintaining and troubleshooting connectivity issues. Considerations include:

- Draining and redirecting traffic flows with Application Load Balancer.
- Managing and troubleshooting VPNs.
- Managing and troubleshooting Cloud Interconnect issues.
- Troubleshooting Cloud Router BGP peering issues.
- Troubleshooting with VPC Flow Logs, firewall logs, and Packet Mirroring.

5.3 Using Network Intelligence Center to monitor and troubleshoot common networking issues. Considerations include:

- Using Network Topology to visualize throughput and traffic flows.
- Using Connectivity Tests to diagnose route and firewall misconfigurations.
- Using Performance Dashboard to identify packet loss and latency (e.g., Google-wide and project scoped).
- Using Firewall Insights to monitor, identify, and improve rules.
- Using Network Analyzer to identify network failures, suboptimal configurations, and utilization warnings.

- Using Flow Analyzer and VPC Flow Logs to evaluate network traffic.

Section 6: Configuring, implementing and managing a cloud network security solution (~13% of the exam)

6.1 Configuring Google Cloud Armor policies. Considerations include:

- Configuring and attaching edge and backend security policies.
- Implementing web application firewall (WAF) rules (e.g., SQL injection, cross-site scripting, and remote file inclusion).
- Configuring advanced network distributed denial of service (DDoS) and Adaptive Protection.
- Configuring rate limiting.
- Configuring bot management.
- Applying Google Threat Intelligence.

6.2 Configuring and managing NGFW policies and VPC Firewall rules. Considerations include:

- Planning the firewall strategy (e.g., VPC firewall rules, Cloud NGFW, hierarchical firewall rules, and third-party integration).
- Understanding the effective policy rules for hierarchical firewall situations.
- Configuring Cloud NGFW to support GKE and Cloud Load Balancing.
- Creating and troubleshooting VPC firewall rules and Cloud NGFW regional/global/hierarchical policies.
- Enabling layer 7 packet inspection with Cloud NGFW Enterprise.
- Migrating from VPC firewall rules to Cloud NGFW policies.
- Configuring VPC and NGFW rule criteria (e.g., rule priority, network protocols, direction [ingress and egress], source, and destination).
- Configuring VPC and Firewall Rules Logging.
- Incorporating micro-segmentation for security purposes (e.g., using metadata, [secure] tags, service accounts, and network tags).
- Differentiating between the different tiers of Cloud NGFW: Essentials, Standard, and Enterprise.

6.3 Configuring and securing internet egress traffic using Public Cloud NAT and Secure Web Proxy. Considerations include:

- Configuring public Cloud NAT IP addressing and assigning automatic and manual Cloud NAT IP addresses.

- Configuring static and dynamic port allocation for Cloud NAT.
- Configuring Secure Web Proxy.

6.4 Configuring self-managed network virtual appliance and Packet Mirroring. Considerations include:

- Routing and inspecting inter-VPC traffic using multi-network interface card (NIC) virtual machines (VMs) (e.g., NGFW appliances).
- Configuring an internal load balancer as a next hop for HA multi-NIC VM routing.
- Configure policy-based routes for HA multi-NIC VM routing.
- Developing a strategy for out-of-band Network Security Integration.
- Configuring Packet Mirroring for VPC traffic toward self-managed collectors.