Google Cloud routes define the paths that network traffic takes from a virtual machine (VM) instance to other destinations. These destinations can be inside your Google Cloud Virtual Private Cloud (VPC) network (for example, in another VM) or outside it.

In a VPC network, a route consists of a single *destination* (CIDR) and a single *next hop*. When an instance in a VPC network sends a packet, Google Cloud delivers the packet to the route's next hop if the packet's destination address is within the route's destination range.

This page provides an overview of how routes work in Google Cloud.

Every VPC network uses a scalable, distributed virtual routing mechanism. Even though some routes can be applied selectively, the routing table (https://wikipedia.org/wiki/Routing_table) for a VPC network is defined at the VPC network level.

Each VM instance has a controller that is kept informed of all applicable routes (#applicable_routes) from the network's routing table. Each packet leaving a VM is delivered to the appropriate next hop of an applicable route based on a routing order. When you add or delete a route, the set of changes is propagated to the VM controllers by using an eventually consistent design (/vpc/docs/using-routes#order_of_operations).

VPC networks can include system-generated routes or custom routes.

The following table summarizes system-generated routes, which are added or updated automatically when you create a VPC network (/vpc/docs/using-vpc#creating_networks), add a subnet (/vpc/docs/using-vpc#add-subnets), expand a subnet's primary IP range (/vpc/docs/using-vpc#expand-subnet), or edit a subnet's secondary IP range (/vpc/docs/using-vpc#edit-secondary). They apply to all instances in the VPC network.

System-generated routes			
Type	**Destination**	**Next hop**	**Removable**

default route (#routingpacketsinternet)	`0.0.0.0/0`		`default-internet-gateway`	Yes
subnet route (#subnet-routes)	Primary and secondary subnet IP ranges (/vpc/docs/vpc#vpc_networks_and_subnets)		VPC network, which forwards packets to VMs in its subnets	Only when the subnet is deleted or when you change the subnet's secondary IP ranges

The following table summarizes custom routes, which you create and maintain directly or by using Cloud Router.

Custom routes

Type	Destination	Next hop	Removable	Applies to
static route (#static_routes)	• IP range that does not partially or exactly overlap with any subnet IP range • IP range *broader* than a subnet IP range	One of: • Instance by name • Instance by IP address • Cloud VPN tunnel	Yes	Either: • All instances in the network • Specific instances in the network, identified by network tag, if the route can be scoped by network tag
dynamic route (#dynamic_routes)	• IP range that does not partially or exactly overlap with any subnet IP range • IP range *broader* than a subnet IP range	IP address of the Cloud Router's BGP peer	Only by a Cloud Router if it no longer receives the route from its BGP peer	• Instances in the same region as the Cloud Router if the VPC network is in regional dynamic routing mode • All instances if the VPC network is in global dynamic routing mode

When you create a VPC network, Google Cloud creates a system-generated default route (https://wikipedia.org/wiki/Default_route). This route serves two purposes:

- It defines the path out of the VPC network, including the path to the internet. In addition to having this route, instances must meet additional requirements (/vpc/docs/vpc#internet_access_reqs) if they need internet access.
- It provides the standard path for Private Google Access (/vpc/docs/private-access-options).

The system-generated default route has a priority of `1000`. Because its destination is the broadest possible (`0.0.0.0/0`), Google Cloud *only* uses it if a route with a more specific destination does *not* apply to a packet. For information about how destination specificity and route priority are used to select a route, see routing order (#routeselection).

If you want to completely isolate your network from the internet or if you need to replace the default route with a custom route, you can delete the default route:

- If you want to route internet traffic to a different next hop, you can replace the default route with a custom static or dynamic route. For example, you could replace it with a custom static route whose next hop is a Cloud VPN tunnel or another instance, such as a proxy server.
- If you remove the default route and do not replace it, packets destined to IP ranges that are not covered by other routes are dropped.

ant: If you don't have custom static routes that meet the routing requirements for Private Google Access
'docs/configure-private-google-access#requirements), deleting the default route might disable Private Google Acces

Subnet routes are system-generated routes that define paths to each subnet in the VPC network.

Each subnet has at least one subnet route whose destination matches the primary IP range of the subnet. If the subnet has secondary IP ranges, Google Cloud creates a subnet route with a corresponding destination for each secondary range. For more information about subnet IP ranges, see networks and subnets (/vpc/docs/vpc#vpc_networks_and_subnets).

No other route can have a destination that matches or is more specific than the destination of a subnet route. You can create a custom route that has a *broader* destination range that *contains* the subnet route's destination range.

In cases of IP range overlap, because Google Cloud uses destination specificity as the first criteria for routing order, the subnet route is always the preferred next hop for packets whose destinations fit in its destination range. This is true even if another route whose destination *contains* the destination of the subnet route has a higher route priority.

Because the subnet route must always have the most specific destination in cases where destination ranges would
, its priority does not matter. The priority of all subnet routes is fixed at **`1000`**.

The following points describe how subnet routes are created and removed:

- When a subnet is created, a corresponding subnet route for the subnet's primary IP range is also created.
    - If you add a secondary IP range to a subnet, a corresponding subnet route for that secondary range is created.
- Auto mode VPC networks (/vpc/docs/vpc#subnet-ranges) create a subnet route for the primary IP ranges of each of their automatically created subnets. You can delete these subnets only if you convert the auto mode VPC network to custom mode.
- You cannot delete a subnet route unless you modify or delete the subnet:
    - When you remove a secondary range from a subnet, the subnet route for that secondary range is deleted automatically.
    - When you delete a subnet, all subnet routes for both primary and secondary ranges are deleted automatically. You cannot delete the subnet route for the subnet's primary range in any other way.
- When networks are connected by using VPC Network Peering (/vpc/docs/vpc-peering), subnet routes from one network are imported into the other network, and vice versa. Thus, all primary and secondary subnet IP ranges must be unique.
    - Subnet routes for subnets in peered networks cannot be removed unless you break the peering relationship. When you break the peering relationship, all imported subnet routes from the other network are automatically removed.

Firewall rules can block communication among instances. For information about instance-to-instance communication, see communication within the network (/vpc/docs/vpc#intra_vpc_reqs).

Custom routes are either static routes that you create manually or dynamic routes maintained automatically by one or more of your Cloud Routers.

Destinations for custom routes *cannot match or be more specific than* any subnet route in the network.

If you are using an auto mode VPC network, don't use destinations that fall into the `10.128.0.0/9` CIDR block because that block defines the current and future address space for subnet routes. For more information, see auto mode IP ranges (/vpc/docs/vpc#ip-ranges).

Static routes with destinations inside `10.128.0.0/9` might be disabled at any time in an auto mode VPC network. This can happen if a new Google Cloud region becomes available and a new subnet is

created automatically (along with its subnet route). For more information, see considerations for auto mode VPC networks (/vpc/docs/vpc#auto-mode-considerations).

Static routes can use any of the static route next hops (#static-route-next-hops). You create static routes in one of two ways:

- You can create them manually (/vpc/docs/using-routes#addingroute).
- If you use the Google Cloud Console to create a Cloud VPN tunnel that uses policy-based routing or one that is a route-based VPN, static routes for the remote traffic selectors are created for you. For more information, see Cloud VPN networks and tunnel routing (/vpn/docs/concepts/choosing-networks-routing).

For more information, see static route parameters (#individualroutes).

Dynamic routes are managed by one or more Cloud Routers. Their destinations always represent IP ranges outside your VPC network, and their next hops are always BGP peer addresses. A Cloud Router can manage dynamic routes for:

- Cloud VPN tunnels that use dynamic routing (/vpn/docs/concepts/choosing-networks-routing#dynamic-routing)
- Cloud Interconnect (/interconnect/docs/how-to/choose-type)

Each instance has a set of *applicable routes*, which are a subset of all routes in the VPC network. Applicable routes are the possible egress paths that a packet can take when sent from the instance.

- **System-generated routes apply to all instances in a VPC network.** The scope of instances to which subnet routes apply cannot be altered; however, you can replace the default route (#routingpacketsinternet).
- **Custom static routes apply to all instances or specific instances.** Static routes with a tag attribute apply to instances that have a corresponding network tag

(/vpc/docs/add-remove-network-tags). If the route doesn't have a network tag, the route applies to all instances in the network.

- When a custom static route has a VM instance next hop, the route is always valid, even if the next-hop VM is deleted, powered off, or malfunctioning. For more information, see instances as next hops (#instance_next_hops).
- When a custom static route has a Cloud VPN tunnel next hop, the route is always valid if the tunnel is up. For information about how Google Cloud handles routes for tunnels when they are down, see when tunnels are down (/vpn/docs/concepts/order-of-routes#tunnels_down) in the Cloud VPN documentation.

- **Dynamic routes apply to instances based on the dynamic routing mode of the VPC network** (/vpc/docs/vpc#routing_for_hybrid_networks)**.** If the VPC network is in regional dynamic routing mode, all Cloud Routers apply routes that they learn within their respective regions. If the network is in global dynamic routing mode, all Cloud Routers apply routes that they learn in the entire network.
  - Cloud Router automatically discards learned custom dynamic routes that correspond to inaccessible next hops (Cloud VPN tunnels that use dynamic routing or Cloud Interconnect). Depending on your network, Cloud Router can take up to 40 seconds of processing time to remove a dynamic route with a next hop that's down.
- **For some load-balanced traffic, the applicable routes originate outside your VPC network.** For more information, see load balancer return paths (#special-lb-paths).

When an instance sends a packet, Google Cloud attempts to select one route from the set of applicable routes according to the following routing order.

Certain load balanced traffic follows a special load balancer return path. This routing order only applies to traffic *not* back through an HTTP(S), SSL Proxy, TCP Proxy, or network load balancer.

1. **Subnet routes are considered first** because Google Cloud requires that subnet routes have the most specific destinations matching the IP address ranges of their respective subnets. Subnet routes are routes for the primary and secondary subnet IP address ranges of each subnet in your VPC network and for the primary and secondary subnet IP address ranges of subnets in peered networks (/vpc/docs/vpc-peering). If the destination for a packet fits inside the destination of a subnet route, it is delivered to that Google Cloud subnet. You *cannot* override a subnet route with any other type of route.

- Google Cloud does not allow you to create a custom static route that has an equal or more specific destination than any subnet route.
- In cases where a static route has the same prefix length as a dynamic route, the static route has a higher priority than the Cloud Router dynamic route. A static route with the same prefix and route metric as a dynamic route always has a higher priority than the dynamic route, so that any conflicting dynamic routes are ignored.
- For custom dynamic routes, Cloud Router ignores any route to other networks received by a Cloud Router if the received route has an equal or more specific destination than any subnet route.
- When you connect VPC networks by using VPC Network Peering, the networks share all subnet routes. Google Cloud prioritizes subnet routes in peer networks similar to the local network's own subnet routes: subnet routes in peer networks must have the most specific destinations.

2. **If the packet does not fit in the destination for a subnet route,** Google Cloud looks for a custom route with the most specific destination.
   - Suppose that the destination for a packet leaving a VM is `10.240.1.4`, and there are two routes with different destinations: `10.240.1.0/24` and `10.240.0.0/16`. Because the most specific destination for `10.240.1.4` is `10.240.1.0/24`, the route whose destination is `10.240.1.0/24` defines the next hop for the packet.
3. **If more than one custom route has the same most specific destination,** Google Cloud uses the following process to select a route from *route candidates*. Route candidates are custom routes (static or dynamic routes) that have the same most specific destination.
   a. If your VPC network is connected to other networks through VPC Network Peering, Google Cloud first narrows the route candidates that are all from a *single* VPC network.

★ **Note:** Google Cloud disregards any route candidates from peer networks if the routes are custom static routes that are scoped to instances by network tag (/vpc/docs/add-remove-network-tags#interaction_with_routes). Network tags from peer networks do not have any meaning outside their local VPC network.

   - If *at least one* route candidate is defined in your *local* VPC network, Google Cloud uses the local route candidate and disregards all candidates from peer networks.
   - If the route candidates are from *multiple* peer networks, Google Cloud selects one of the peer networks where at least one candidate route is defined and disregards candidate routes from other peer networks. Google Cloud uses a non-deterministic

method to select the single peer network, regardless of each route's priority. If you add or remove networks from the peering group, Google Cloud might choose candidate routes from a different peer network.

Google Cloud moves on to select a route from a single network's candidate routes.

b. If a route with the highest priority is available, Google Cloud sends the packet to its next hop.

c. If multiple routes have the highest priority, Google Cloud selects a route based on whether it's a static or dynamic route and the value of its next hop. Google Cloud uses the following order (most to least preferred):

   i. Google Cloud selects a custom static route whose next hop is Next Hop Instance, Next Hop IP, or Next Hop VPN Tunnel.

   ii. Google Cloud selects a custom dynamic route that is being advertised by a Cloud Router.

   iii. Google Cloud selects a custom static route whose next hop is Next Hop Internal TCP/UDP Load Balancer.

   iv. Google Cloud selects a custom static route by using Next Hop Default Internet Gateway.

d. If no single route can be selected, Google Cloud distributes traffic among the next hops of the route candidates by using a five-tuple hash for affinity, implementing an ECMP (https://en.wikipedia.org/wiki/Equal-cost_multi-path_routing) routing design. The hash is calculated from the protocol, the source and destination IP addresses, and the source and destination ports of the packet being sent. This calculation is done for each packet sent based on the number of route candidates that are available. If the number of available route candidates changes, the hash might direct the packet to a different next hop.

4. **If no applicable destination is found,** Google Cloud drops the packet, replying with an ICMP destination or network unreachable error.

tant: Google Cloud does *not* consider geographical distance when sending traffic to an applicable route. Keep this in
create custom static routes with next hops of other instances or Cloud VPN tunnels. Instances are zonal resources, a
unnels are regional resources. If the route is applicable and is an appropriate next hop according to the routing order,
ces in other zones and regions can use it, no matter how far away the next hop is geographically.

If you use one of the following Google Cloud load balancers, Google Cloud has special routes for the load balancers and their associated health checks, which are described in more detail in the following sections. These routes are defined *outside* your VPC network, in Google's production network. They don't appear in your VPC network's routing table. You cannot remove or override them, even if you delete or replace a default route in your VPC network.

For these load balancer types, Google Front End (GFE) systems serve as proxies. When a client sends a request to the load balancer, a proxy terminates the TCP session and opens a new TCP session with your backend VM. Routes defined outside your VPC network facilitate communication from GFE proxies to your backend VMs and from your backend VMs to GFE proxies.

For more information, see the following pages:

- HTTP(S) Load Balancing (/load-balancing/docs/https/)
- SSL Proxy Load Balancing (/load-balancing/docs/ssl/)
- TCP Proxy Load Balancing (/load-balancing/docs/tcp/)

When a client on the internet sends a TCP or UDP request through a network load balancer to a backend VM, Google Cloud uses routes defined outside your VPC network to facilitate communication from the client to your backend VM and from your backend VM to the client.

For more information, see Network Load Balancing (/load-balancing/docs/network/).

Packets sent from Google Cloud health check probe systems have sources as described in Probe IP ranges and firewall rules (/load-balancing/docs/health-check-concepts#ip-ranges). Routes that facilitate communication between Google Cloud health check probe systems and your backend VMs exist outside your VPC network, and cannot be removed. However, your VPC network must have ingress allow firewall rules to permit traffic from these systems.

Each static route consists of the following components:

- **Name** and **Description.** These fields identify the route. A name is required, but a description is optional. Every route in your *project* must have a unique name.
- **Network.** Each route must be associated with exactly one VPC network.
- **Destination range.** The destination range is a single IPv4 CIDR block that contains the IP addresses of systems that receive incoming packets. Google Cloud does not support IPv6 destination ranges. Destinations must be expressed in CIDR notation (https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing), and the broadest destination possible is `0.0.0.0/0`.
- **Priority.** If multiple routes have identical destinations, priority is used to determine which route should be used. Lower numbers indicate *higher* priorities; for example, a route with a priority value of `100` has higher priority than one with a priority value of `200`.
- **Next hop.** Static routes can have next hops that point to the default internet gateway, a Google Cloud instance, or a Cloud VPN tunnel. For more information, see static route next hops (#static-route-next-hops).
- **Tags.** You can specify a list of network tags (/vpc/docs/add-remove-network-tags) so that the route *only* applies to instances that have at least one of the listed tags. If you don't specify tags, Google Cloud applies the route to all instances in the network.

The following are valid next hops for static routes. For more information about each type, see the `gcloud` reference documentation (/sdk/gcloud/reference/compute/routes/create).

- **Next Hop Gateway (`next-hop-gateway`).** You can specify a default internet gateway to define a path to public IP addresses.
- **Next Hop Instance (`next-hop-instance`).** You can direct traffic to *an existing instance* in Google Cloud by specifying its *name* and *zone*. Traffic is directed to the primary internal IP address of the VM's network interface in the VPC network where you define the route.
  - If the primary internal IP address for the VM instance's network interface in the VPC network changes, Google Cloud updates the VPC network's routing table automatically so that traffic continues to be sent to the VM at its new IP address.
  - If the VM is replaced by a new VM with the same name in the same zone, Google Cloud updates the VPC network's routing table automatically so that traffic is directed to the replacement VM.

- Google Cloud only validates that a next hop VM exists *when you create the route*. It does not validate that the VM is able to process or forward packets. If the VM is deleted but *not* replaced, the route *still* applies, which might result in dropped packets. For more information, see instances as next hops (#instance_next_hops).

- **Next Hop Internal TCP/UDP Load Balancer (`next-hop-ilb`).** For Internal TCP/UDP Load Balancing, you can use a load balancer's IP address as a next hop that distributes traffic among healthy backend instances. For example, you can use a custom static route whose next hop is an internal TCP/UDP load balancer to distribute traffic among backend VMs. Custom static routes that use this next hop cannot be scoped to specific instances by network tag.

- **Next Hop IP (`next-hop-address`).** You can provide an IP address as a next hop if that address fits within either the primary or secondary IP range of an existing subnet in your VPC network. You cannot use a public IP address or an internal IP address in an on-premises network.

  - When you use `next-hop-address`, Google Cloud passes traffic to any VM instance assigned to that IP address. If you replace a VM instance with another one that uses the same IP address, traffic goes to the replacement, even if it has a different name. To define a next hop by VM name, use Next Hop Instance instead.

  - Google Cloud only checks that the IP address is a valid member of a subnet's primary or secondary IP range when you *create the route*. It does not validate that an instance exists at the Next Hop IP address. If the Next Hop IP address is not assigned to any instance, the route *still* applies, which might result in dropped packets. For more information, see Considerations for instance and load balancer next hops (#instance_next_hops).

- **Next Hop VPN Tunnel (`next-hop-vpn-tunnel`).** For Cloud VPN tunnels that use policy-based routing and route-based VPNs (/vpn/docs/concepts/choosing-networks-routing#ts-tun-routing), you can direct traffic to the VPN tunnel by creating routes whose next hops refer to the tunnel by its name and region. Google Cloud ignores routes whose next hops are Cloud VPN tunnels that are down. For more examples about how routes and Cloud VPN tunnels interact, see the Cloud VPN routing examples (/vpn/docs/concepts/order-of-routes#routing-examples).

Custom static routes cannot use target instances (/sdk/gcloud/reference/compute/target-instances/) as next hops

Instance-based routing refers to a static route with a next hop that is a VM instance (`next-hop-instance` or `next-hop-address`).

Internal TCP/UDP load balancer as a next hop refers to a static route with a next hop that is an internal TCP/UDP load balancer (`next-hop-ilb`).

When you configure instance-based routing or an internal TCP/UDP load balancer as a next hop, consider the following guidelines:

- **You must configure the backend VMs or the Next Hop Instance to forward packets from other sources (`can-ip-forward`)** (/vpc/docs/using-routes#canipforward)**.** You can enable IP forwarding on a per-VM basis when you create the VM. To enable IP forwarding for VMs created automatically as part of a managed instance group, you must enable IP forwarding in the instance template used by the instance group. You must make this configuration change *in addition* to any operating system configuration necessary to route packets.
- **Software running on the backend VM or the Next Hop Instance must be configured appropriately.** For example, third-party appliance VMs that act as routers or firewalls must be configured according to the manufacturer's instructions.
- **Backend VMs or the Next Hop Instance must have appropriate firewall rules.** You must configure firewall rules that apply to the packets being routed. Keep the following in mind:
    - Ingress firewall rules applicable to instances that perform routing functions must include the IP addresses of *routed packet sources*. The implied deny ingress rule blocks all incoming packets, so you must at least create custom ingress allow firewall rules.
    - Egress firewall rules applicable to instances that perform routing functions must include the IP addresses of *routed packet destinations*. The implied allow egress rule permits this unless you've created a specific egress deny rule to override it.
    - Take into account whether the backend VM or Next Hop Instance is performing Network Address Translation (NAT) when creating your firewall rules.

  For more information, see Implied firewall rules (/vpc/docs/firewalls#default_firewall_rules).

- When you use an instance as a next hop (`next-hop-instance` or `next-hop-address`), remember that the instance is a zonal resource. Selecting a zone implies that you have selected a region. Google Cloud does not consider regional distance for routes that use a Next Hop Instance, so it is possible to create a route that sends traffic to a next hop in a different region. This adds egress costs and introduces network latency.
- Google Cloud performs one of the following basic validations *only when you create the route*:
    - If you specify `next-hop-instance`, the instance must already exist.

  - If you specify `next-hop-address`, the IP address must be in an existing subnet's primary or secondary IP range.
- If you, for example, remove an instance or delete a subnet later, Google Cloud still considers any route that uses that instance as a next hop when it evaluates applicable routes (#instancerouting) . This might cause some or all packets for a given destination to be dropped, depending on the other routes in your network.
- If software in a VM instance (such as the VM's operating system or the VM's process that routes packets) fails, packets destined for that instance are dropped. To enhance reliability, consider using an internal TCP/UDP load balancer as a next hop (/load-balancing/docs/internal/ilb-next-hop-overview) instead. An internal TCP/UDP load balancer requires a configured health check, and this health check determines how new connections are routed (/load-balancing/docs/internal/#traffic_distribution).
- Disabling a network interface by configuring the guest operating system of the instance does not cause Google Cloud to stop considering the interface to be a route's next hop.

- To create and manage routes, see Using routes (/vpc/docs/using-routes).
- To get an overview of Google Cloud VPC networks, see VPC network overview (/vpc/docs/vpc).
- To create, modify, or delete VPC networks, see Using VPC networks (/vpc/docs/using-vpc).