Private services access is a private connection between your VPC network and a network owned by Google or a third party. Google or the third party, entities who are offering services, are also known as *service producers*. The private connection enables VM instances in your VPC network and the services that you access to communicate exclusively by using internal (RFC 1918 (https://tools.ietf.org/html/rfc1918)) IP addresses. VM instances don't need Internet access or external IP addresses to reach services through private services access.

To learn more about private services access and other private access options, see Private Access Options for Services (/vpc/docs/private-access-options).

At a high level, to use private services access, you must allocate (#allocating-range) an IP address range (CIDR block) in your VPC network and then create a private connection (#creating-connection) to a service producer.

To establish a private connection, complete the following prerequisites:

- Check that the service you're using supports (/vpc/docs/private-access-options#private-services-supported-services) private services access.
- You must have an existing VPC network that you will use to connect to the service producer's network. VM instances must use this VPC network to connect to services over a private connection.
- Activate (/service-infrastructure/docs/service-networking/getting-started) the Service Networking API (/service-infrastructure/docs/service-networking/reference/rest/) in your project. The API is required to create a private connection.
- Create a Cloud project or choose an existing one. To learn how to create a Cloud project, see Creating and Managing Projects (/resource-manager/docs/creating-managing-projects).
- Install the Cloud SDK (/sdk/docs) if you want to run the `gcloud` command-line examples in this guide.

Project owners, editors, and IAM members with the Network Admin (/compute/docs/access/iam#compute.networkAdmin) role can create allocated IP address ranges and manage private connections.

For more information on roles, read the VPC IAM roles (/compute/docs/access/iam) documentation.

If you are using Shared VPC (/vpc/docs/shared-vpc), create the allocated IP range and private connection in the host project. Typically, a network administrator in the host project must do these tasks. After the host project is set up, VM instances in service projects can use the private connection.

Because a private connection is implemented as a VPC peering connection, the same quota and limits (/vpc/docs/using-vpc-peering#limits) that apply to VPC Network Peering also apply to private services access.

Before you create a private connection, you must allocate an IP address range to be used by the service producer's VPC network. This ensures that there's no IP address collision between your VPC network and the service producer's network. Create an allocated range for each service producer.

When you allocate a range in your VPC network, that range is ineligible for subnets (primary and secondary ranges) and destinations of custom static routes.

Ensure that your IP address range allocations don't overlap with dynamic routes. Google Cloud doesn't check if your
:ion overlap with dynamic routes (routes learned through BGP, such as through Cloud Router).

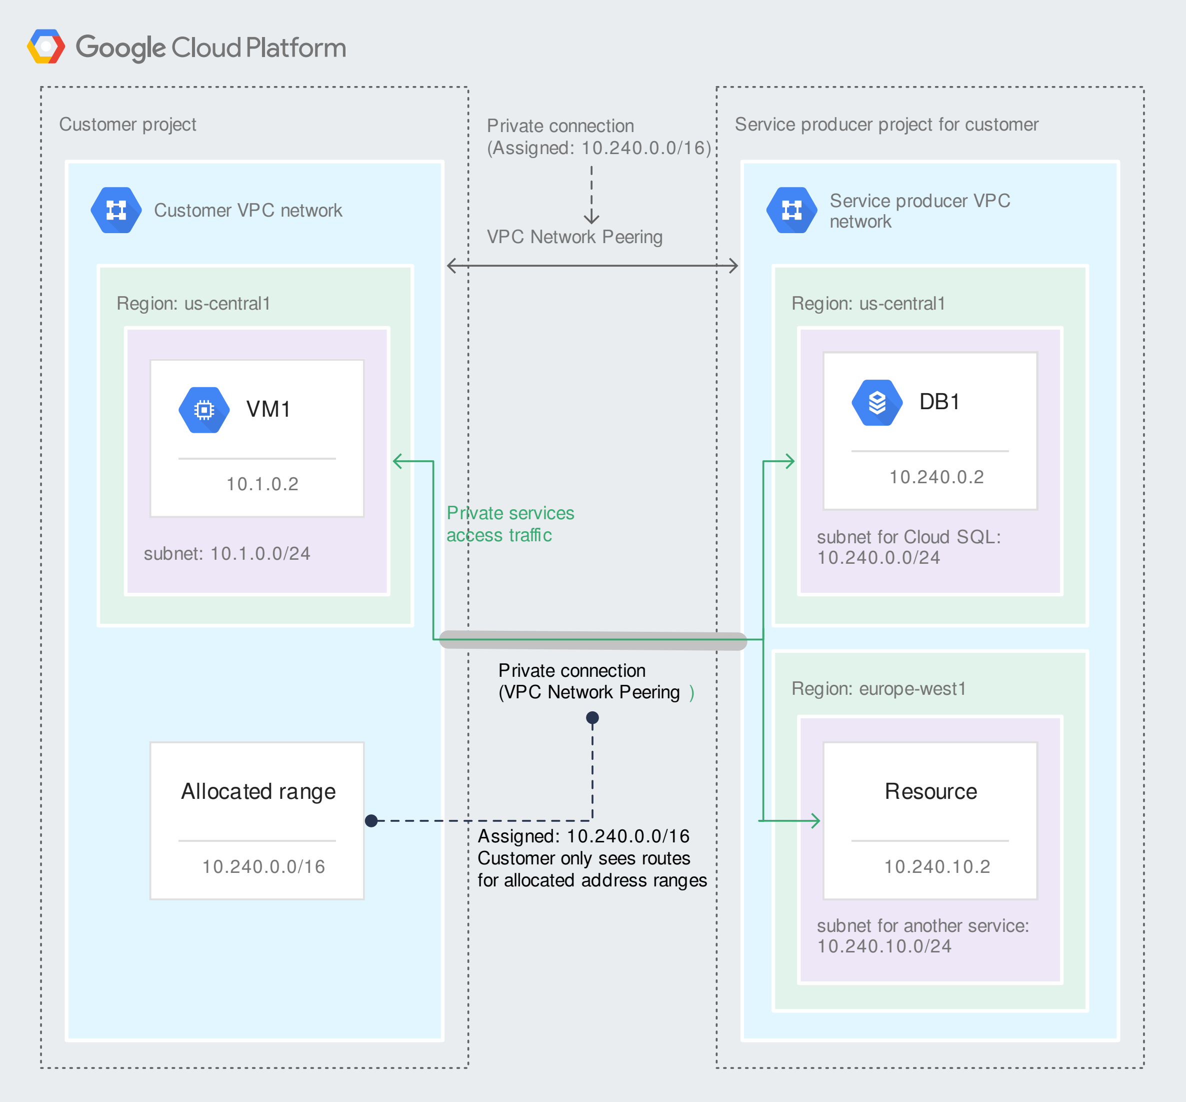

(/vpc/images/private-services-access-ranges.svg)
Private services access ranges (click to enlarge)

When a service producer creates a subnet on their side of the connection, an open range from the allocation is selected for the subnet's IP address range.

Each service producer requires a minimum IP address range size. For Google, the minimum size is a single `/24` block (256 addresses), but the recommended size is a `/16` block (65,536 addresses). The size you choose depends on factors such as the number of services and regions you use. You might use a single service in multiple regions for redundancy scenarios or to reduce latency.

For example, if you use two services in three different regions, the service producer must create 6 subnets, each with a `/24` block. If you plan to use additional services or regions, you might quickly exhaust a `/20` allocation. If don't have a contiguous `/16` block, you can start with a smaller allocation and add new ones (#modify-ip-range) if you need more IP addresses later.

When you establish a private connection and create a resource with a private IP address, the service creates a subnet in which to provision the resource. The service selects an available IP address range from the allocated range. You cannot select or modify the service producer's subnet IP address range. Even if you delete the private connection or the allocated IP address range, the subnet remains until you delete all resources in the subnet.

As you provision additional resources, the service provisions them in existing regional subnets that it previously created. If a subnet is full, the service creates a new one in that region.

Before you allocate an IP address range, consider the following constraints:

- Select a range that doesn't overlap with existing allocated ranges, subnets, or custom static routes. No two ranges can overlap.
- If you're using an auto mode (/vpc/docs/vpc#subnet-ranges) VPC network, you can't create an allocated range that matches or overlaps with `10.128.0.0/9`. This range is for automatically created subnets (/vpc/docs/vpc#ip-ranges).
- Select a CIDR block that is large enough to meet your current and future needs. If you later find that the range isn't sufficient in size, expand (#modify-ip-range) the range if possible. Although you can assign multiple allocations to a single service producer, Google enforces a quota on the number of IP address ranges that you can allocate but not the size (netmask) of each range.
- Don't reuse the same allocated range for multiple service producers. Although it's possible, doing so can lead to IP address overlap. Each service producer has visibility only into their network and can't know which IP addresses other service producers are using.
- You can only assign one CIDR block to an allocated range when you create the allocation. If you need to expand (#modify-ip-range) the IP address range, you can't add more blocks to an allocation. Instead, you can create another allocation or recreate the existing one by using a larger block that encompasses the new and existing ranges.
- If you create the allocation yourself instead of having Google do it (such as through Cloud SQL), you can use the same naming convention to signal to other users or Google services that an allocation for Google already exists. When a Google service allocates a range on your

behalf, the service uses the following format to name the allocation: `google-managed-services-[your network name]`. If this allocation exists, Google services use the existing one instead of creating another one.

The following steps describe how to create an allocated IP address range.

You can list ranges with the `--filter` flag to see which ranges you can use for private services access. Filter for ranges with the purpose `VPC_PEERING`, as shown in the following example:

After you create an allocated range, you can create a private connection to a service producer. The private connection establishes a VPC Network Peering (/vpc/docs/vpc-peering) connection between your VPC network and the service producer's network.

Private connections are a one-to-one relationship between your VPC network and a service producer. If a single service producer offers multiple services, you only need one private connection for all of the producer's services.

If you connect to multiple service producers, use a unique allocation for each service producer. This practice helps you manage your network settings, such as routes and firewall rules, for each service producer.

If you have an on-premises network connected to your VPC, you can configure the peering connection so that on-pre
can communicate with the service producer's network. For more information, refer to the on-premises host (#on-pre
eshooting topic.

After you create a private connection, you can list it to check that it exists. The list also shows the list of allocated ranges that are associated with each connection. For example, if you don't remember which allocated range you assigned to a connection, view the list to find out.

For existing private connections, you can add or remove allocated IP address ranges without disrupting traffic. For example, as you scale, you might add an allocated range if you're close to exhausting the existing one.

To delete a private connection, you must delete the corresponding VPC peering connection. Your VPC network is disconnected from the service producer's VPC network, and existing resources in both networks remain but lose private services access. You can re-establish connectivity by creating a private connection (#creating-connection) again.

Before you delete an allocated IP address range, check (#listing-connections) that no private connection is using it. You can delete (#removing-connection) or modify (#modifying-conntection) an existing private connection to disassociate the allocated range. If you don't, existing connections remain active, but there's nothing preventing your VPC from using IP addresses that overlap with the service producer's network. Also, the service can't create new subnets because there's no allocated IP address range to select from.

When you create a private connection with a service producer, you allocate an IP address range for them to use. If you use multiple services from a service producer, each service will reserve a chunk of IP addresses from that allocated range. You can check which services are using which IP addresses so that, for example, you can see which services are using large blocks of IP addresses and avoid IP address exhaustion.

To view which service is using a particular IP address range:

1. List your private connections (#listing-connections).
2. Find the peering connection name that connects you to the relevant service producer.
3. List the routes (/vpc/docs/using-routes#listingroutes) for your VPC network.
4. Find the routes with a next hop that match the peering connection name. The destination range of the routes indicates which IP addresses each service is using.

For a given private connection, if you exhaust your allocated IP address space, you can expand the existing allocation or add new ones. The expanded allocation must be a contiguous IP address range that includes the existing range. Expanding an allocation is recommended because there's no limit on the size of an allocation, but there is a limit on the number of allocations that you can create.

To expand an existing allocation:

1. List your private connections (#listing-connections) and record the name of the allocated range you need to expand.
2. Delete (#deleting-allocation) the existing allocated range.
3. Create a new allocated range (#allocating-range) by using the same name as the deleted range. Specify an IP address range that includes the deleted IP address range. That way, existing peered resources that are using the old allocated range can continue to use the same IP

addresses without colliding with resources in your VPC network. For example, if the previous allocated range was `192.168.0.0/20`, create a new allocated range as `192.168.0.0/16`.

To add allocated ranges to an existing private connection:

1. Create a new allocated range (#allocating-range). This range doesn't have to be contiguous with existing allocated ranges.
2. Add (#modifying-connection) the allocated range to the exiting private connection.

:ant: Don't delete an allocated range that is used by a private connection unless you need to replace it with a new, larg
:ed range. Don't replace an allocated range that's in use with a smaller range. If you delete or shrink an allocated rang
n use, you prevent the service producer from being able to create new subnets on your behalf in their network.

The service producer's network might not have the correct routes to direct traffic to your on-premises network. By default, the service producer's network only learns the subnet routes from your VPC network. Therefore, any request that's not from a subnet IP range is dropped by the service producer.

In your VPC network, update the peering connection (/vpc/docs/using-vpc-peering#update-peer-connection) to export custom routes to the service producer's network. Exporting routes sends all eligible static and dynamic routes (/vpc/docs/vpc-peering#considerations) that are in your VPC network, such as routes to your on-premises network, to the service producer's network. The service producer's network automatically imports them and then can send traffic back to your on-premises network through the VPC network.