Each TPU node runs on a peer VPC network (/vpc/docs/vpc-peering) that Google manages. When you create a TPU node, the system automatically creates a peer connection to a VPC network (/vpc/docs/vpc) from your project. The peer connection provides access to the TPU node from your instances, GKE clusters, and other Cloud Platform services that run on the same VPC network.

When you specify an internal IP address range for your TPU node, you can use any valid address in your VPC internal network as long as it does not conflict with other addresses that are already used on your network.

The internal TPU IP address is only accessible by members of your project.

If you are creating a TPU node in a smaller project on a network that you are not sharing with other project members, you can use address such as `192.168.0.0` or `172.16.0.0` and the range size is specified automatically for you.

Alternatively, the `ctpu` (/tpu/docs/ctpu-reference) utility can automatically configure a range when you use it to create a TPU node.

For TPU nodes in larger automated production environments where many systems share addresses on your internal VPC network, select an address range that you know will be available for larger numbers of hosts in your TPU node.

Cloud TPU **VPC Service Controls** can provide additional security for your Cloud TPU resources. To learn more about V e Controls, see the VPC Service Controls overview (/vpc-service-controls/docs/overview). To learn about the limitatio Cloud TPU with VPC Service Controls, see the supported products and limitations service-controls/docs/supported-products).

IP address ranges for TPU nodes must meet the following requirements:

- The IP address range must be from within the internal IP address ranges:

- Your range must be one of the following formats:

  where:

  - `W,X,Y,Z` are integers in the range 0-255 with no leading zeros.
  - `/N` is an optional address range size. If you do not specify a range size Cloud TPU selects the correct range size for your TPU type.
- Your range must have enough addresses to accommodate the size of your TPU type. If you omit the range size, Cloud TPU selects the correct range size for your TPU type. If you must specify a range size, select a range size that provides enough addresses for 1/4th the number of TPU cores in your node with at least 8 addresses at a minimum. For example, if you select a `v3-512` TPU type with 512 cores, your range must have a size of `/25` to provide 128 addresses for the TPU peer network. For both `v3-8` and `v3-32` TPU types, you must specify the minimum range size of `/29` to provide at least 8 addresses.
- You must select a range that does not conflict with another network resource on the network that you are using. For example, if you deploy your Cloud TPU on the "default" network that is created on standard Compute Engine projects, then this network will already have a subnetwork for `us-central1` using the `10.128.0.0/20` range and you cannot place any of your subnetworks inside this range. This is true for all other regions.

When you are writing applications to automatically create and manage TPU nodes, you can reserve allocated IP address ranges (/vpc/docs/configure-private-services-access#allocating-range) to automatically identify open IP address ranges on your network and reserve them to be used for your TPU nodes.

If you are manually creating and deleting TPU nodes for specific jobs, the **`ctpu`** (/tpu/docs/ctpu-reference) utility automatically configures IP address ranges for you.