# Concepts

## Cloud TPU

### Cloud TPUs
An overview of Cloud TPU resources.
(https://cloud.google.com/tpu/docs/tpus)

### Services that can access TPUs
Helps you decide which Cloud TPU service (Compute Engine, Google Kubernetes Engine, or AI Platform) best suits your needs.
(https://cloud.google.com/tpu/docs/deciding-tpu-service)

### TPU types and zones
Details the types of Cloud TPU and their zones of availability.
(https://cloud.google.com/tpu/docs/types-zones)

### Internal IP address ranges
Setting up internal IP addresses for Cloud TPU nodes.
(https://cloud.google.com/tpu/docs/performance-guide)

### Performance guide
High-level performance tips for running TensorFlow programs on Cloud TPU.
(https://cloud.google.com/tpu/docs/performance-guide)

### Training on TPU Pods
An overview of training on Cloud TPU Pods.

(https://cloud.google.com/tpu/docs/training-on-tpu-pods)

### Using TPU Estimator API

The differences between training a TensorFlow model with the standard TensorFlow Estimator and with the TPUEstimator API.

(https://cloud.google.com/tpu/docs/using-estimator-api)

### Processing large images with TPUs

Using spatial partitioning with Cloud TPUs and TPUEstimator for training with very large images and video.

(https://cloud.google.com/tpu/docs/spatial-partitioning)

### System Architecture

An overview of Cloud TPU hardware and software architecture.

(https://cloud.google.com/tpu/docs/system-architecture)

### Using bfloat16

Using the bfloat16 format to store activations and gradients in memory.

(https://cloud.google.com/tpu/docs/bfloat16)

### Advanced Guide to Inception v3

An advanced view of the guide to running Inception v3 on Cloud TPU.

(https://cloud.google.com/tpu/docs/inception-v3-advanced)