

Restore from a snapshot

Archived: 2026-04-06 00:16:36 UTC

[Skip to main content](#)

- [Technology areas](#)
 - [Overview](#)
 - [Guides](#)
 - [APIs & Reference](#)
 - [Samples](#)
 - [Resources](#)
- [Cross-product tools](#)
- [Console](#)

- Discover
 - [Product overview](#)
 - [Compute Engine instances](#)
 - [Instance groups](#)
 -
 -

- Get started
 -
 -

- Create instances
 - [Instance creation overview](#)
 -
 -
 -

- Create multiple VMs
 -
 -
 -

- Create sole-tenant VMs
 - [Sole-tenancy overview](#)
 - [Create sole-tenant node templates](#)

- [Create sole-tenant node groups](#)
- [Provision a sole-tenant VM](#)
- [Advanced maintenance control for sole-tenant nodes](#)
- [Sole-tenancy best practices](#)
- [Sole-tenancy accounting FAQ](#)
-
-
-
-
- Migrate VMs
- [Choose a migration path](#)
- [Bring your own licenses](#)
-
-
- [Move an existing VM to a new VM](#)
- Connect to VMs
- Connect to a VM
 - [About SSH connections](#)
 - Linux VMs
 - [Connect to VMs](#)
 -
 - [Connect as the root user](#)
 - [Connect using service accounts](#)
 - [Configure apps to use SSH](#)
 -
 - Windows VMs
 - [Connect to Windows VMs using RDP](#)
 - [Connect to a Windows VM's SAC](#)
 - [Connect to Windows VMs using SSH](#)
 - [Connect to Windows VMs using PowerShell](#)
- Manage access to VMs
 - Linux VMs
 - [Choose an access management method](#)
 - [About OS Login](#)
 - [Set up OS Login](#)
 - [Set up OS Login to require SSH certificates](#)

- [Enable security keys with OS Login](#)
- [Manage OS Login in an organization](#)
- [Monitor OS Login audit logs](#)
-
-
-
- [Manage tags for resources](#)
-
- [IP addresses](#)
-
- [Create a PTR record for a VM](#)
- [Verify VM identity](#)
- Manage storage
- [Choose a disk type](#)
-
-
- Add disks to VMs
 - [Create a VM with Local SSD disks](#)
 - [Create a VM with additional non-boot disks](#)
 - [Create a new Hyperdisk](#)
 - [Create a new Persistent Disk](#)
 - [Add disks from a storage pool to VMs](#)
 - [Share a disk between VMs](#)
 - [Attach a disk to a VM](#)
 - [Mount in-memory RAM disks](#)
-
- [View disk details](#)
-
-
-
-
- Make disks highly available
 -
 -
- Back up and restore
- [Data protection options](#)
- [Configure the default backup setting](#)
- Back up VMs

-
-
-
-
- [Recover a VM with a corrupted or full disk](#)
- Manage VMs
- Basic operations and lifecycle
 - [VM instance lifecycle](#)
 -
 - View VM properties
 - [Detect if a VM is running in Compute Engine](#)
 - [Get a list of VMs](#)
 - [Get the details of a VM](#)
 - [Get the UUID of a VM](#)
 - [View the source image of a VM](#)
 - [View referrers to VMs](#)
 - [View network configuration of an instance](#)
 - [View the number of visible CPU cores](#)
 -
 -
 -
- [Update VM tenancy](#)
- Manage multiple VMs
 -
 -
- Host maintenance events
 - [About host events](#)
 - [Live migration process](#)
 - [Set the host maintenance policy](#)
 - [Query metadata server for notices](#)
 - [Simulate a host maintenance event](#)
 - [Handle GPU host maintenance events](#)
 - [Monitor and plan for a host maintenance event](#)
 - [Manually start host maintenance](#)
-
- Securing VMs

- [About Shielded VMs](#)
- [About Confidential VMs](#)
- [Protect resources with VPC Service Controls](#)
- [Monitor security risks with Security Command Center](#)
- Manage operating systems
 -
 - [Manage operating systems using VM Manager](#)
- Manage OS images
 - [Image management best practices](#)
 - [Image families best practices](#)
 - [Access Red Hat Knowledgebase](#)
 - [Manage access to custom images](#)
 - [Set up trusted image policies](#)
 - [Export a custom image to Cloud Storage](#)
 - [Set image versions in an image family](#)
 - [Deprecate a custom image](#)
 - [Delete a custom image](#)
- [Manage OS packages](#)
- Manage licenses
 - [About licenses](#)
 - [Manage licenses](#)
 - [License changes and restrictions](#)
 - [Switch between PAYG and BYOS](#)
 - [Switch Windows Server from BYOL to PAYG](#)
 - [Update licenses after an OS upgrade](#)
 - [Append RHEL ELS licenses](#)
 - [Upgrade from Ubuntu to Ubuntu Pro](#)
-
-
- [Run shutdown scripts](#)
-
- [Enable the virtual random number generator \(Virtio RNG\)](#)
- Deploy workloads
 -
 - [Agent for Compute Workloads overview](#)
 -
- Applications

- [Interactive: Build a to-do app with MongoDB](#)
- [Deploy an ASP.NET application](#)
- [Set up Joomla](#)
- [Set up LAMP](#)
- [Perform blue/green deployments using Cloud Build](#)
-
- Databases
 - MySQL
 - [MySQL on Compute Engine](#)
 - [Install MySQL on Compute Engine](#)
 - [Configure MySQL on Compute Engine](#)
 - [Set up client access with a private IP address](#)
 - [Cloning a MySQL database on Compute Engine](#)
 - [Architectures for high availability of MySQL clusters on Compute Engine](#)
 - [Deploying a highly available MySQL 5.6 cluster with DRBD on Compute Engine](#)
 -
 - SQL Server
 - [Best practices for SQL Server VMs](#)
 - Create
 - [Create a high-performance SQL Server VM](#)
 - [Add a SQL Server license to an existing Linux server](#)
 - [Add a SQL Server license to an existing Windows server](#)
 -
 - Configure
 - [Set up AlwaysOn availability groups using an internal load balancer](#)
 - [Set up AlwaysOn availability groups using a distributed network name](#)
 - [Set up a failover cluster VM that uses S2D](#)
 - [Set up a failover cluster VM with multi-writer disks](#)
 - [Set up a SQL Server cluster on Linux with Always On availability groups and Pacemaker](#)
 - Migrate
 - [Migrate a SQL Server database from AWS EC2 to Compute Engine](#)
 - [Migrate a SQL Server database from Windows to Linux](#)
 - Disaster recovery

- [Disaster recovery for Microsoft SQL Server](#)
- [Disaster recovery for Microsoft SQL server on Persistent disk](#)
- [Disaster recovery for Microsoft SQL server on Hyperdisk](#)
- [Deploying Microsoft SQL Server for multi-regional disaster recovery](#)
- [Backup SQL Server databases to a Google Cloud Storage bucket](#)
- [Cloning a Microsoft SQL Server database on Compute Engine](#)
- [Load test SQL Server using HammerDB](#)
-
- Containers
 - [Containers on Compute Engine](#)
 - [Deploy containers on VMs and managed instance groups](#)
 - [Configure options to run your container](#)
 -
 - OpenShift workloads
 - [OpenShift on Google Cloud overview](#)
 - [Plan for OpenShift on Google Cloud](#)
 - [Overview of Cluster Services for OpenShift](#)
 - [Built-in integrations for OpenShift](#)
 - [Best practices for high availability with OpenShift](#)
 - [Disaster recovery for OpenShift on Google Cloud](#)
 - [Disaster recovery strategies for active-passive and active-inactive setups with OpenShift](#)
- Microsoft Windows
 - [Windows workloads](#)
 - [Best practices for Windows Server VMs](#)
 - [Setting up Active Directory](#)
 - [Best practices for running Active Directory on Google Cloud](#)
 - [Deploy Microsoft SharePoint Server on Compute Engine](#)
 - [Deploying Microsoft Exchange Server 2016 on Compute Engine](#)
 -
-
- Others
 -
 -
 -
- Monitor
- Monitor logs

- [View audit logs](#)
- [View usage reports](#)
- [View Compute Engine operations](#)
- [Migrate from activity logs to audit logs](#)
- [View activity logs](#)
-
- [Organize resources using labels](#)
- Scale
-
- [Autoscale node groups](#)
-
- Load balancing
 - [About load balancing and scaling](#)
 - [Add an instance group to a load balancer](#)
 - [Request routing to a multi-region external HTTPS load balancer](#)
 - [Cross-region load balancing for Microsoft IIS backends](#)
 - [Set up Internal TCP/UDP Load Balancing](#)
-
- Optimize
-
-
-
-
-
- Network performance
 - [Network bandwidth](#)
 - [Use Google Virtual NIC](#)
 - [Use IRDMA network driver](#)
 - [Use IDPF network interface](#)
 - [Configure a VM with higher bandwidth](#)
 - [Reduce latency by using compact placement policies](#)
 -
 - [Benchmark higher bandwidth VMs](#)
 - [Optimize app latency with load balancing](#)
 - [Use DPDK to improve network performance](#)
 -
- Troubleshoot

- [General tips](#)
-
- Troubleshoot VMs
 -
 -
 -
 -
 -
-
-
-
-
-
-
-
-
-
-
-

Restore from a snapshot Stay organized with collections Save and categorize content based on your preferences.

- On this page
- [Before you begin](#)
 - [Required roles](#)
- [Limitations](#)
- [Create a disk from a snapshot and optionally attach it to an instance](#)
 - [Create an instance from existing disks](#)
- [Create an instance from a boot disk snapshot](#)
- [Create a compute instance from a non-boot disk snapshot](#)
-

A disk is either a boot disk that is used to start and run the operating system on a compute instance or a non-boot disk that an instance uses only for data storage.

You can use snapshots to backup and restore disk data in the following ways:

- After you take a snapshot of a boot or non-boot disk, [create a new disk based on the snapshot](#).
- After you take a snapshot of a boot disk, [create a new instance based on the boot disk snapshot](#).
- After you take a snapshot of a non-boot disk, [create a new instance with a new non-boot disk based on the snapshot](#).

Before you begin

- If you haven't already, set up [authentication](#). Authentication verifies your identity for access to Google Cloud services and APIs. To run code or samples from a local development environment, you can authenticate to Compute Engine by selecting one of the following options:

Select the tab for how you plan to use the samples on this page:

When you use the Google Cloud console to access Google Cloud services and APIs, you don't need to set up authentication.

1. [Install](#) the Google Cloud CLI. After installation, [initialize](#) the Google Cloud CLI by running the following command:

```
gcloud init
```

If you're using an external identity provider (IdP), you must first [sign in to the gcloud CLI with your federated identity](#).

2. [Set a default region and zone](#).

To use the Go samples on this page in a local development environment, install and initialize the gcloud CLI, and then set up Application Default Credentials with your user credentials.

1. [Install](#) the Google Cloud CLI.
2. If you're using an external identity provider (IdP), you must first [sign in to the gcloud CLI with your federated identity](#).
3. If you're using a local shell, then create local authentication credentials for your user account:

```
gcloud auth application-default login
```

You don't need to do this if you're using Cloud Shell.

If an authentication error is returned, and you are using an external identity provider (IdP), confirm that you have [signed in to the gcloud CLI with your federated identity](#).

For more information, see [Set up authentication for a local development environment](#).

To use the Java samples on this page in a local development environment, install and initialize the gcloud CLI, and then set up Application Default Credentials with your user credentials.

1. [Install](#) the Google Cloud CLI.
2. If you're using an external identity provider (IdP), you must first [sign in to the gcloud CLI with your federated identity](#).
3. If you're using a local shell, then create local authentication credentials for your user account:

```
gcloud auth application-default login
```

You don't need to do this if you're using Cloud Shell.

If an authentication error is returned, and you are using an external identity provider (IdP), confirm that you have [signed in to the gcloud CLI with your federated identity](#).

For more information, see [Set up authentication for a local development environment](#).

To use the Node.js samples on this page in a local development environment, install and initialize the gcloud CLI, and then set up Application Default Credentials with your user credentials.

1. [Install](#) the Google Cloud CLI.
2. If you're using an external identity provider (IdP), you must first [sign in to the gcloud CLI with your federated identity](#).
3. If you're using a local shell, then create local authentication credentials for your user account:

```
gcloud auth application-default login
```

You don't need to do this if you're using Cloud Shell.

If an authentication error is returned, and you are using an external identity provider (IdP), confirm that you have [signed in to the gcloud CLI with your federated identity](#).

For more information, see [Set up authentication for a local development environment](#).

To use the Python samples on this page in a local development environment, install and initialize the gcloud CLI, and then set up Application Default Credentials with your user credentials.

1. [Install](#) the Google Cloud CLI.
2. If you're using an external identity provider (IdP), you must first [sign in to the gcloud CLI with your federated identity](#).
3. If you're using a local shell, then create local authentication credentials for your user account:

```
gcloud auth application-default login
```

You don't need to do this if you're using Cloud Shell.

If an authentication error is returned, and you are using an external identity provider (IdP), confirm that you have [signed in to the gcloud CLI with your federated identity](#).

For more information, see [Set up authentication for a local development environment](#).

To use the REST API samples on this page in a local development environment, you use the credentials you provide to the gcloud CLI.

[Install](#) the Google Cloud CLI.

If you're using an external identity provider (IdP), you must first [sign in to the gcloud CLI with your federated identity](#).

For more information, see [Authenticate for using REST](#) in the Google Cloud authentication documentation.

Required roles

To get the permissions that you need to restore from a snapshot, ask your administrator to grant you the [Compute Instance Admin \(v1\)](#)(`roles/compute.instanceAdmin.v1`) IAM role on the project. For more information about granting roles, see [Manage access to projects, folders, and organizations](#).

This predefined role contains the permissions required to restore from a snapshot. To see the exact permissions that are required, expand the **Required permissions** section:

Required permissions

The following permissions are required to restore from a snapshot:

- To create a disk from a globally scoped snapshot (default):
 - `compute.disks.create` on the project
 - `compute.instances.attachDisk` on the instance
 - `compute.disks.use` on the disk to attach
 - `compute.snapshots.useReadOnly` , `compute.snapshots.create` , or `compute.disks.createSnapshot` on the project
- To create a disk from a regionally scoped snapshot:
 - `compute.disks.create` on the project
 - `compute.instances.attachDisk` on the instance
 - `compute.disks.use` on the disk to attach
 - `compute.snapshots.useReadOnly` on the project
- To create an instance from a boot disk and non-boot disk snapshot, at a minimum, you need the following permissions:
 - `compute.instances.create` on the project
 - `compute.snapshots.useReadOnly` on the snapshot
 - `compute.disks.create` on the project
 - `compute.disks.use` on the disk
 - `compute.instances.attachDisk` on the instance

You might also be able to get these permissions with [custom roles](#) or other [predefined roles](#).

Limitations

- The new disk must be at least the same size as the original source disk for the snapshot. If you create a disk that is larger than the original source disk for the snapshot, you must [resize the file system on that disk](#) to include the additional disk space. Depending on your operating system and file system type, you might need to use a different file system resizing tool. For more information, see your operating system documentation.
- You can create a new zonal or regional disks from a given snapshot at most once every ten minutes. If you want to issue a burst of requests to snapshot your disks, you can issue at most 6 requests in 60 minutes. This limit does not apply when creating regional disks from a snapshot. For more information, see [Snapshot frequency limits](#).

Create a disk from a snapshot and optionally attach it to an instance

If you backed up a boot or non-boot disk with a snapshot, you can create a new disk based on the snapshot.

1. In the Google Cloud console, go to the **Snapshots** page.

[Go to Snapshots](#)

2. Find the name of the snapshot that you want to restore.

3. Go to the **Disks** page.

[Go to the Disks page](#)

4. Click **Create new disk**.

5. Specify the following configuration parameters:

- A name for the disk.
- A type for the disk.
- Optionally, you can override the default region and zone selection. You can select any region and zone, regardless of the storage location of the source snapshot.

6. Under **Source type**, click **Snapshot**.

7. Select the name of the snapshot to restore.

8. Select the size of the new disk, in gigabytes. This number must be equal to or larger than the original source disk for the snapshot.

9. Click **Create** to create the disk.

Optionally, you can then [attach a non-boot disk to a instance](#).

1. Use the `gcloud compute snapshots list` [command](#) to find the name of the snapshot you want to restore:

```
gcloud compute snapshots list
```

2. Use the [gcloud compute snapshots describe command](#) to find the size of the snapshot you want to restore:

```
gcloud compute snapshots describe SNAPSHOT_NAME
```

Replace *SNAPSHOT_NAME* with the name of the snapshot being restored.

3. Use the [gcloud compute disks create command](#) to create a new regional or zonal disk from your snapshot. You can include the `--type` flag to specify the type of disk to create.

- To create a zonal disk from a globally scoped snapshot:

```
gcloud compute disks create DISK_NAME \  
  --zone=ZONE \  
  --size=DISK_SIZE \  
  --source-snapshot=SNAPSHOT_NAME \  
  --type=DISK_TYPE
```

- ([Preview](#)) To create a zonal disk from a regionally scoped snapshot:

```
gcloud beta compute disks create DISK_NAME \  
  --zone=ZONE \  
  --source-snapshot=SNAPSHOT_NAME \  
  --source-snapshot-region=SOURCE_REGION \  
  --type=DISK_TYPE
```

- To create a regional disk from a globally scoped snapshot:

```
gcloud beta compute disks create DISK_NAME \  
  --size=DISK_SIZE \  
  --source-snapshot=SNAPSHOT_NAME \  
  --type=DISK_TYPE \  
  --region=REGION \  
  --replica-zones=ZONE1,ZONE2
```

- ([Preview](#)) To create a regional disk from a regionally scoped snapshot:

```
gcloud beta compute disks create DISK_NAME \  
  --size=DISK_SIZE \  
  --source-snapshot=SNAPSHOT_NAME \  
  --source-snapshot-region=SOURCE_REGION \  
  --type=DISK_TYPE
```

```
--region=REGION \  
--replica-zones=ZONE1,ZONE2
```

Replace the following:

- *DISK_NAME*: the name of the new disk
- *DISK_SIZE*: the size of the new disk, in gibibytes (GiB). This number must be equal to or larger than the original source disk for the snapshot.
- *SNAPSHOT_NAME*: the name of the snapshot being restored
- *DISK_TYPE*: the [type](#) of the disk, for example, `pd-ssd`, `hyperdisk-throughput` or `hyperdisk-balanced-high-availability`
- *REGION*: the region for the regional disk to reside in, for example: `europe-west1`
- *SOURCE_REGION*: the region that the source snapshot is scoped to
- *ZONE*: the zone where the new disk will reside
- *ZONE1,ZONE2*: the zones within the region where the two disk replicas are located, for example: `europe-west1-b` and `europe-west1-c`

4. Optional: Attach the new disk to an existing instance by using the [gcloud compute instances attach-disk command](#):

```
gcloud compute instances attach-disk INSTANCE_NAME \  
--disk DISK_NAME
```

Replace the following:

- *INSTANCE_NAME*: the name of the instance
- *DISK_NAME*: the name of the disk made from your snapshot

Go

Before trying this sample, follow the Go setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Go API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Java

Before trying this sample, follow the Java setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Java API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Node.js

Before trying this sample, follow the Node.js setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Node.js API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Python

Before trying this sample, follow the Python setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Python API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

1. Construct a `GET` request to `snapshots.list` to display the list of snapshots in your project.

```
GET https://compute.googleapis.com/compute/v1/projects/PROJECT_ID/global/snapshots
```

Replace *PROJECT_ID* with your project ID.

2. Construct a `POST` request to create a zonal disk or a [regional](#) disk using the respective `disks.insert` method:
 - For zonal disks: `disks.insert`
 - For regional disks: `regionDisks.insert`

Include the `name`, `sizeGb`, and `type` properties. To restore a disk using a snapshot, you must include the `sourceSnapshot` property.

- To create a zonal disk from a globally scoped snapshot:

```
POST https://compute.googleapis.com/compute/v1/projects/PROJECT_ID/zones/ZONE/disks
{
  "name": "DISK_NAME",
  "sourceSnapshot": "SNAPSHOT_NAME",
  "sizeGb": "DISK_SIZE",
  "type": "zones/ZONE/diskTypes/DISK_TYPE"
}
```

- ([Preview](#)) To create a zonal disk from a regionally scoped snapshot:

```
POST https://compute.googleapis.com/compute/beta/projects/PROJECT_ID/zones/ZONE/disks
{
  "name": "DISK_NAME",
  "sourceSnapshot": "projects/PROJECT_ID/regions/SOURCE_REGION/snapshots/SNAPSHOT_NAME",
  "sizeGb": "DISK_SIZE",
}
```

```

"type": "projects/PROJECT_ID/zones/ZONE/diskTypes/DISK_TYPE",
"zone": "projects/PROJECT_ID/zones/ZONE"
}

```

- To create a regional disk from a globally scoped snapshot:

```

POST https://compute.googleapis.com/compute/beta/projects/PROJECT_ID/regions/REGION/disks
{
  "name": "DISK_NAME",
  "sourceSnapshot": "SNAPSHOT_NAME",
  "region": "projects/PROJECT_ID/regions/REGION",
  "replicaZones": [
    "projects/PROJECT_ID/zones/ZONE1",
    "projects/PROJECT_ID/zones/ZONE2"
  ],
  "sizeGb": "DISK_SIZE",
  "type": "zones/ZONE/diskTypes/DISK_TYPE"
}

```

- ([Preview](#)) To create a regional disk from a regionally scoped snapshot:

```

POST https://compute.googleapis.com/compute/beta/projects/PROJECT_ID/regions/REGION/disks
{
  "name": "DISK_NAME",
  "sourceSnapshot": "projects/PROJECT_ID/regions/SOURCE_REGION/snapshots/SNAPSHOT_NAME",
  "replicaZones": [
    "projects/PROJECT_ID/zones/ZONE1",
    "projects/PROJECT_ID/zones/ZONE2"
  ],
  "sizeGb": "DISK_SIZE",
  "type": "projects/PROJECT_ID/regions/REGION/diskTypes/DISK_TYPE"
}

```

Replace the following:

- *PROJECT_ID*: your project ID
- *ZONE*: the zone where your instance and new disk are located
- *DISK_NAME*: the name of the new disk
- *SNAPSHOT_NAME*: the source snapshot for the disk that you're restoring
- *REGION*: the region for the regional disk to reside in, for example: `europa-west1`
- *SOURCE_REGION*: the region that the source snapshot is scoped to
- *ZONE1*, *ZONE2*: the zones where replicas of the new disk should be located
- *DISK_SIZE*: the size of the new disk, in gibibytes (GiB). This number must be equal to or larger than the original source disk for the snapshot.

- *DISK_TYPE*: full or partial URL for the [type](#) of the disk, for example, `PROJECT_ID/zones/ZONE/diskTypes/pd-ssd` , `PROJECT_ID/zones/ZONE/diskTypes/hyperdisk-balanced` or `PROJECT_ID/zones/ZONE/diskTypes/hyperdisk-balanced-high-availability`

3. Optional. Attach the new disk to an existing instance.

Construct a `POST` request to the [instances.attachDisk](#) [method](#), and include the URL to the disk that you just created from your snapshot.

For zonal disks:

```
POST https://compute.googleapis.com/compute/v1/projects/PROJECT_ID/zones/ZONE/instances/INSTAN
{
  "source": "/compute/v1/projects/PROJECT_ID/zones/ZONE/disks/DISK_NAME"
}
```

For regional disks:

```
POST https://compute.googleapis.com/compute/v1/projects/PROJECT_ID/zones/ZONE/instances/INSTAN
{
  "source": "/compute/v1/projects/PROJECT_ID/regions/REGION/disks/DISK_NAME"
}
```

Replace the following:

- *PROJECT_ID*: your project ID
- *ZONE*: the zone where your instance and new disk are located
- *REGION*: the region where the regional disk is located. This must be the same region that the compute instance is located in.
- *INSTANCE_NAME*: the name of the instance where you're adding the new disk
- *DISK_NAME*: the name of the new disk

After you create and attach a new disk to an instance, you must [mount the disk](#) so that the operating system can use the available storage space.

Create an instance from existing disks

You can create boot disks and data disks from snapshots and then attach these disks to a new compute instance.

Go

Before trying this sample, follow the Go setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Go API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Java

Before trying this sample, follow the Java setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Java API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Node.js

Before trying this sample, follow the Node.js setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Node.js API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Python

Before trying this sample, follow the Python setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Python API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Create an instance from a boot disk snapshot

If you created a snapshot of the boot disk of a compute instance, you can use that snapshot to create a new instance.

To quickly create more than one instance with the same boot disk, [create a custom image](#), then [create instances from that image](#) instead of using a snapshot.

To create a compute instance with a regional boot disk that was created from a snapshot, use the Google Cloud CLI or REST.

1. In the Google Cloud console, go to the **Create an instance** page.

[Go to Create an instance](#)

If prompted, select your project and click **Continue**. The **Create an instance** page appears and displays the **Machine configuration** pane.

2. In the **Machine configuration** pane, do the following:
 1. In the **Name** field, specify a name for your instance. For more information, see [Resource naming convention](#).
 2. Optional: In the **Zone** field, select a zone for this instance.

The default selection is **Any**. If you don't change this default selection, then Google automatically chooses a zone for you based on machine type and availability.

3. Select the machine family for your instance. The Google Cloud console then displays the machine series that are available for your selected machine family. The following machine family options are available:

- **General purpose**
- **Compute optimized**
- **Memory optimized**
- **Storage optimized**
- **GPUs**

4. In the **Series** column, select the machine series for your instance.

If you selected **GPUs** as the machine family in the previous step, then select the **GPU type** that you want. The machine series is then automatically selected for the selected GPU type.

5. In the **Machine type** section, select the machine type for your instance.

3. In the navigation menu, click **OS and storage**. In the **Operating system and storage** pane that appears, configure your boot disk by doing the following:

1. Click **Change**. The **Boot disk** pane appears and displays the **Public images** tab.
2. Click **Snapshots**. The **Snapshot** tab appears.
3. In the **Snapshot** list, select the snapshot to use.
4. In the **Boot disk type** list, select the type of the boot disk.
5. In the **Size (GB)** field, specify the size of the boot disk.
6. Optional: For advanced configuration options, expand the **Show advanced configurations** section.
7. To confirm your boot disk options and return to the **Operating system and storage** pane, click **Select**.

4. In the navigation menu, click **Networking**. In the **Networking** pane that appears, do the following:

1. Go to the **Firewall** section.
2. To permit HTTP or HTTPS traffic to the instance, select **Allow HTTP traffic** or **Allow HTTPS traffic**.

The Google Cloud console adds a network tag to your instance and creates the corresponding ingress firewall rule that allows all incoming traffic on `tcp:80` (HTTP) or `tcp:443` (HTTPS). The network tag associates the firewall rule with the instance. For more information, see [Firewall rules overview](#) in the Virtual Private Cloud documentation.

5. Optional: Specify other configuration options. For more information, see [Configuration options during instance creation](#).

6. To create and start the instance, click **Create**.

Use the `gcloud compute instances create` [command](#) and include the `--source-snapshot` flag.

```
gcloud compute instances create INSTANCE_NAME

  --source-snapshot=BOOT_SNAPSHOT_NAME

  --boot-disk-size=BOOT_DISK_SIZE

  --boot-disk-type=BOOT_DISK_TYPE

  --boot-disk-device-name=BOOT_DISK_NAME
```

Replace the following:

- `INSTANCE_NAME` : [name](#) for the new instance
- `BOOT_SNAPSHOT_NAME` : name of the boot disk snapshot that you want to restore to the boot disk of the new instance
- `BOOT_DISK_SIZE` : Optional: size, in GiB, of the new boot disk

The size must be equal to or larger than the size of the source disk from which the snapshot was made.

- `BOOT_DISK_TYPE` : Optional: [type](#) of the boot disk, for example `PROJECT_ID/zones/ZONE/diskTypes/pd-ssd` or `PROJECT_ID/zones/ZONE/diskTypes/hyperdisk-balanced`
- `BOOT_DISK_NAME` : name of the new boot disk for this instance

Use the `gcloud compute instances create` [command](#) and include the `--create-disk` flag with the `source-snapshot` , `replica-zones` , and `boot` properties.

```
gcloud compute instances create INSTANCE_NAME

  --zone=ZONE

  --create-disk=^:^name=DISK_NAME:source-snapshot=BOOT_SNAPSHOT_NAME:boot=true:type=BOOT_DISK_TY
```

The characters `^:^` specify that a colon `:` is used as the separator between each of the disk properties. This is required so that you can use a comma `,` when specifying the zones for `replica-zones` .

Replace the following:

- `INSTANCE_NAME` : [name](#) for the new instance
- `ZONE` : To zone to create the instance in
- `DISK_NAME` : Optional: a name for the disk

- `BOOT_SNAPSHOT_NAME` : name of the boot disk snapshot that you want to restore to the boot disk of the new instance.
- `BOOT_DISK_TYPE` : Optional: [type](#) of the boot disk, for example `pd-ssd` or `hyperdisk-balanced-high-availability`
- `REMOTE_ZONE` : The region that the boot disk is replicated to. The `replica-zones` property requires two zones separated by comma, and one of the zones must be the same as the zone for the instance.

Go

Before trying this sample, follow the Go setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Go API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Java

Before trying this sample, follow the Java setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Java API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Node.js

Before trying this sample, follow the Node.js setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Node.js API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

Python

Before trying this sample, follow the Python setup instructions in the [Compute Engine quickstart using client libraries](#). For more information, see the [Compute Engine Python API reference documentation](#).

To authenticate to Compute Engine, set up Application Default Credentials. For more information, see [Set up authentication for client libraries](#).

When you use the API to create an instance from a snapshot, the following restrictions apply:

- Only one disk can be used as the boot disk.
- You must attach the boot disk as the first disk for that instance.
- If you specify the `source` property, you cannot also specify the `initializeParams` property. Providing a `source` indicates that the boot disk exists already, but the `initializeParams` property indicates that Compute Engine should create a new boot disk.

To create an instance from a boot disk snapshot, use the [instances.insert method](#) and specify the `sourceSnapshot` field under the `disks` property. You can optionally specify the `diskSizeGb` and `diskType` properties for the new boot disk.

```
POST https://compute.googleapis.com/compute/v1/projects/PROJECT_ID/zones/ZONE/instances
{
  "name": "INSTANCE_NAME",
  "machineType": "machineTypes/MACHINE_TYPE",
  "disks": [{
    "boot": true,
    "initializeParams": {
      "sourceSnapshot": "global/snapshots/BOOT_SNAPSHOT_NAME",
      "diskSizeGb": "BOOT_DISK_SIZE",
      "diskType": "BOOT_DISK_TYPE"
    }
  }],
  "networkInterfaces": [
    {
      "nicType": "GVNIC"
    }
  ]
}
```

Replace the following:

- `PROJECT_ID` : your project ID
- `ZONE` : zone where you want to create the new instance
- `INSTANCE_NAME` : [name](#) of the instance that you want to restore a snapshot to
- `MACHINE_TYPE` : [machine type](#) of the instance
- `BOOT_SNAPSHOT_NAME` : name of the snapshot that you want to use to create the boot disk of the new instance
- `BOOT_DISK_SIZE` : Optional: size, in gibibytes (GiB), for the new boot disk

The size must be equal to or larger than the size of the source disk from which the snapshot was made.

- `BOOT_DISK_TYPE` : Optional: [type](#) of the boot disk, for example `PROJECT_ID/zones/ZONE/diskTypes/pd-ssd` or `PROJECT_ID/zones/ZONE/diskTypes/hyperdisk-balanced`

To create a compute instance with a regional boot disk using a boot disk snapshot as the source, use the [instances.insert method](#) and specify the `sourceSnapshot` and `replicaZones` fields in the `disks` property.

```
POST https://compute.googleapis.com/compute/v1/projects/PROJECT_ID/zones/ZONE/instances
{
  "name": "INSTANCE_NAME",
  "disks": [
```

```

{
  "boot": true,
  "initializeParams": {
    "sourceSnapshot": "global/snapshots/BOOT_SNAPSHOT_NAME",
    "replicaZones": [
      "projects/PROJECT_ID/zones/ZONE",
      "projects/PROJECT_ID/zones/REMOTE_ZONE"
    ],
    "diskType": "BOOT_DISK_TYPE"
  }
},
],
"networkInterfaces": [
  {
    "nicType": "GVNIC"
  }
]
}

```

Replace the following:

- *PROJECT_ID* : your project ID
- *ZONE* : the name of the zone where you want to create the instance
- *INSTANCE_NAME* : a name for the instance
- *BOOT_SNAPSHOT_NAME* : the name of the boot disk snapshot
- *REMOTE_ZONE* : the remote zone for the regional disk
- *BOOT_DISK_TYPE* : Optional: [type](#) of the boot disk, for example `PROJECT_ID/zones/ZONE/diskTypes/pd-ssd` or `PROJECT_ID/zones/ZONE/diskTypes/hyperdisk-balanced-high-availability`

Create a compute instance from a non-boot disk snapshot

If you backed up a non-boot disk with a snapshot, you can create an instance with an additional non-boot disk based on the snapshot.

When restoring non-boot snapshots to a new instance from the console, first create a disk from each snapshot. Then, attach the new disks when you create the instance.

1. Restore each non-boot snapshot to a new disk.
 1. In the Google Cloud console, go to the **Disks** page.

[Go to Disks](#)

2. Click **Create disk**.

3. Specify a **Name** for your disk. For more information, see [Resource naming convention](#).
4. Select the **Region** and **Zone** for this disk. The disk and instance must be in the same zone for zonal disks, or region for regional disks.
5. Select a disk **Type**.
6. Under **Source type**, select **Snapshot**.
7. Under the new **Source snapshot** field, select a non-boot snapshot that you want to restore to the new disk.
8. To create the disk, click **Create**.

Repeat these steps to create a disk from each snapshot that you want to restore.

2. In the Google Cloud console, go to the **VM instances** page.

[Go to VM instances](#)

3. Select your project and click **Continue**.
4. Click **Create instance**.
5. Specify a **Name** for your instance. For more information, see [Resource naming convention](#).
6. Select the **Region** and **Zone** for this instance. The disk and instance must be in the same zone for zonal disks, or region for regional disks.
7. Select a **Machine type** for your instance.
8. If you want to allow incoming external traffic, change the **Firewall** rules for the instance.
9. To attach disks to the instance, expand the **Advanced options** section, and then do the following:
 1. Expand the **Disks** section.
 2. Click **Attach existing disk**.
 3. In the **Disk** list, select a disk to attach to this instance.
 4. In the **Attachment Setting** section, select disk's attachment **Mode** and the **Deletion rule**. For more information about adding new disks, see [Add a Persistent Disk](#) or [Add Hyperdisk](#).
 5. Click **Save**.

Repeat these steps for each disk that you want to attach.

10. To create and start the instance, click **Create**.

Create an instance by using the `gcloud compute instances create` [command](#). For each non-boot snapshot that you want to restore, include the `--create-disk` flag, and specify a `source-snapshot` .

For example, to restore two snapshots of non-boot disks to a new instance, use the following command:

```
gcloud compute instances create INSTANCE_NAME \
  --create-disk source-snapshot=SNAPSHOT_1_NAME,name=DISK_1_NAME,size=DISK_1_SIZE,type=DISK_1_TYPE \
  --create-disk source-snapshot=SNAPSHOT_2_NAME,name=DISK_2_NAME,size=DISK_2_SIZE,type=DISK_2_TYPE
```

Replace the following:

- *INSTANCE_NAME* : [name](#) for the new instance
- *SNAPSHOT_1_NAME* and *SNAPSHOT_2_NAME* : names of the non-boot disk snapshots that you want to restore
- *DISK_1_NAME* and *DISK_2_NAME* : names of the new non-boot disks to create for this instance
- *DISK_1_SIZE* and *DISK_2_SIZE* : Optional: sizes, in gibibytes (GiB), of each new non-boot disk

The sizes must be equal to or larger than the sizes of the source disks from which the snapshot was made.

- *DISK_1_TYPE* and *DISK_2_TYPE* : Optional: the disk [types](#) to create, for example `pd-ssd` or `hyperdisk-balanced`

When using REST to restore a non-boot snapshot to a new instance, the following restrictions apply:

- Only one disk can be the boot disk.
- You must attach the boot disk as the first disk for that instance.
- If you specify the `source` property, you can't also specify the `initializeParams` property. Providing a `source` indicates that the boot disk exists already, but the `initializeParams` property indicates that Compute Engine should create a new boot disk.

Create a `POST` request to the [instances.insert method](#) and specify the `sourceSnapshot` field under the `initializeParams` property. You can add multiple non-boot disks by repeating the `initializeParams` property for every non-boot disk that you want to create. You can optionally specify the `diskSizeGb` and `diskType` properties for any of the disks that you create.

For example, to restore two non-boot disk snapshots to a new instance, make the following request:

```
POST https://compute.googleapis.com/compute/v1/projects/PROJECT_ID/zones/ZONE/instances
{
  "name": "INSTANCE_NAME",
  "machineType": "machineTypes/MACHINE_TYPE",
  "networkInterfaces": [
    {
      "nicType": "GVNIC"
    }
  ],
  "disks": [
    {
      "autoDelete": "true",
      "boot": "true",
      "diskSizeGb": "BOOT_DISK_SIZE",
```

```

"diskType": "BOOT_DISK_TYPE",
"initializeParams": {
  "sourceImage": "projects/IMAGE_PROJECT/global/images/family/IMAGE_FAMILY"
}
},
{
"deviceName": "DEVICE_1_NAME",
"initializeParams": {
  "sourceSnapshot": "global/snapshots/SNAPSHOT_1_NAME",
  "diskSizeGb": "DISK_1_SIZE",
  "diskType": "DISK_1_TYPE"
}
},
{
"deviceName": "DEVICE_2_NAME",
"initializeParams": {
  "sourceSnapshot": "global/snapshots/SNAPSHOT_2_NAME",
  "diskSizeGb": "DISK_2_SIZE",
  "diskType": "DISK_2_TYPE"
}
}
]
}

```

Replace the following:

- `PROJECT_ID` : your project ID
- `ZONE` : zone where you want to create the instance
- `INSTANCE_NAME` : a [name](#) for the new instance
- `MACHINE_TYPE` : [machine type](#) of the instance
- `DISK_SIZE` : Optional: size, in gibibytes (GiB), of the corresponding disk

When provided, this property must be equal to or larger than the size of the source disk from which the snapshot was made.

- `DISK_TYPE` : Optional: full or partial URL for the [type](#) of the corresponding disk, for example, `PROJECT_ID/zones/ZONE/diskTypes/pd-ssd` or `PROJECT_ID/zones/ZONE/diskTypes/hyperdisk-balanced`
- `IMAGE_PROJECT` : the project containing the image. For example, `debian-cloud`
- `IMAGE_FAMILY` : an image family. This creates the instance from the most recent, non-deprecated OS image in that family. For example, if you specify `"sourceImage": "projects/debian-cloud/global/images/family/debian-11"`, the Compute Engine creates an instance using the latest version of the OS image in the Debian 11 image family.
- `DEVICE_NAME` : Optional: the [device name](#) displayed in the guest OS of the instance

- `SNAPSHOT_NAME` : the names of corresponding non-boot disk snapshots that you want to restore to new disks on the instance

Except as otherwise noted, the content of this page is licensed under the [Creative Commons Attribution 4.0 License](#), and code samples are licensed under the [Apache 2.0 License](#). For details, see the [Google Developers Site Policies](#). Java is a registered trademark of Oracle and/or its affiliates.

Last updated 2026-04-02 UTC.

Source: <https://cloud.google.com/compute/docs/disks/restore-and-delete-snapshots>