

Creating instance templates

This page describes how to create and manage [instance templates](/compute/docs/instance-templates) (/compute/docs/instance-templates). Instance templates let you define the machine type, boot disk image or container image, network, and other instance properties to use when creating virtual machine (VM) instances. You can use instance templates to create instances in a [managed instance group \(MIG\)](/compute/docs/instance-groups) (/compute/docs/instance-groups) or to [create individual instances](/compute/docs/instances/create-vm-from-instance-template) (/compute/docs/instances/create-vm-from-instance-template).

Before you begin

- If you want to use the command-line examples in this guide:
 1. Install or update to the latest version of the [gcloud command-line tool](/compute/docs/gcloud-compute) (/compute/docs/gcloud-compute).
 2. [Set a default region and zone](/compute/docs/gcloud-compute#set_default_zone_and_region_in_your_local_client) (/compute/docs/gcloud-compute#set_default_zone_and_region_in_your_local_client).
- If you want to use the API examples in this guide, [set up API access](/compute/docs/api/prereqs) (/compute/docs/api/prereqs).

Creating a new instance template

Most of the instance properties that you can define in a regular API request to create an individual VM instance can be defined in the instance template, including any instance metadata, startup scripts, persistent disks, service accounts, and so on.

At a minimum, the same required properties for creating an instance are also required to create an instance template. For a list of the required fields, see the [instanceTemplates\(\).insert](/compute/docs/reference/latest/instanceTemplates/insert) (/compute/docs/reference/latest/instanceTemplates/insert) reference.

Create an instance template through the [Google Cloud Console](https://console.cloud.google.com/) (https://console.cloud.google.com/), [gcloud command-line tool](/compute/docs/gcloud-compute) (/compute/docs/gcloud-compute), or the [API](/compute/docs/reference/latest/instanceTemplates) (/compute/docs/reference/latest/instanceTemplates).

1. In the Cloud Console, go to the **Instance templates** page.

[Go to the Instance templates page](https://console.cloud.google.com/compute/instanceTemplates) (<https://console.cloud.google.com/compute/instanceTemplates>)

2. Click **Create instance template**.
3. Fill in the fields you want for your instance template, or accept the default values. The following values are provided by default:

- Machine type: `n1-standard-1`
- Image: The latest Debian image
- Boot disk: A new standard boot disk named after the instance
- VPC network: The `default` VPC network
- IP address: An ephemeral external IP address

4. Optionally, change the instance's **Shielded VM** (</security/shielded-cloud/shielded-vm>) settings, if you chose an image that supports Shielded VM:

- a. Click the **Security** tab in the **Management, Security, Disks, Networking, Sole Tenancy** section.
- b. If you want to disable Secure Boot, clear the **Turn on Secure Boot** checkbox. Secure Boot helps protect your VM instances against boot-level and kernel-level malware and rootkits. For more information, see [Secure boot](/security/shielded-cloud/shielded-vm#secure-boot) (</security/shielded-cloud/shielded-vm#secure-boot>).
- c. If you want to disable the virtual trusted platform module (vTPM), clear the **Turn on vTPM** checkbox. The vTPM enables Measured boot, which validates the VM pre-boot and boot integrity. For more information, see [Virtual Trusted Platform Module \(vTPM\)](/security/shielded-cloud/shielded-vm#vtpm) (</security/shielded-cloud/shielded-vm#vtpm>).

★ **Important:** Disabling the vTPM also disables integrity monitoring because integrity monitoring relies on data gathered by Measured boot.

- d. If you want to disable integrity monitoring, uncheck the **Turn on Integrity Monitoring** checkbox. Integrity monitoring lets you monitor the boot integrity of your Shielded VM instances by using Cloud Monitoring. For more information, see [Integrity monitoring](/security/shielded-cloud/shielded-vm#integrity-monitoring) (</security/shielded-cloud/shielded-vm#integrity-monitoring>).
5. Optionally, click the tabs in the **Management, Security, Disks, Networking, Sole Tenancy** section to further customize your template. For example, you can add up to 15 secondary non-boot disks.

6. Optionally, click **Equivalent REST** at the bottom of the page to view the REST request body, which includes the [JSON representation](#) (/compute/docs/reference/rest/v1/instanceTemplates) of your instance template.
7. Click **Create** to create the template.

Creating an instance template based on an existing instance

You can use the Compute Engine API or `gcloud` tool to save the configuration of an existing VM instance as an instance template.

You can override how the source disks are defined in the template.

If you need to override other properties, first create an instance template based on an existing instance, then [create a similar template](#) (#based-on-existing-template) with additional overrides.

Warning: If the existing instance contains a static external IP address, that address is copied into the instance template and might limit the use of the template.

`gcloudAPI` (#api)

For a list of all the flags, see the `gcloud instance-templates create` (/sdk/gcloud/reference/compute/instance-templates/create) command.

```
gcloud compute instance-templates create [INSTANCE_TEMPLATE_NAME] \
  --source-instance=[SOURCE_INSTANCE] \
  --source-instance-zone=[SOURCE_INSTANCE_ZONE] \
  [--configure-disk= \
    device-name=[SOURCE_DISK], \
    instantiate-from=[INSTANTIATE_FROM], \
    auto-delete=[AUTO_DELETE]]
```

where:

- `[INSTANCE_TEMPLATE_NAME]` is the name of the template to create.
- `[SOURCE_INSTANCE]` is the name of the instance to use as a model for the new template.

- `[SOURCE_INSTANCE_ZONE]` is the zone that contains the source instance.
- `[SOURCE_DISK]` is the name of a source-instance disk that you want to override within the template.
- `[INSTANTIATE_FROM]` specifies whether to include the disk and which image to use. Valid values depend on the type of disk:
 - `source-image` or `source-image-family` (valid only for boot and other persistent read/write disks).
 - `custom-image` (valid only for boot and other persistent read/write disks). If specified, then the path or URL for the custom image must also be specified, as shown in the following example.
 - `attach-read-only` (valid only for read-only disks).
 - `blank` (valid only for non-boot persistent disks and local SSDs). If specified, then, when the template is used to create a new instance, the disk is created unformatted. You must [format and mount](/compute/docs/disks/add-persistent-disk#formatting) (`/compute/docs/disks/add-persistent-disk#formatting`) the disk in a startup script before you can use it in a scalable setup.
 - `do-not-include` (valid only for non-boot persistent disks and read-only disks).
 - `blank` (valid only for non-boot persistent disks and read-only disks).
- `[AUTO_DELETE]` specifies whether the disk is auto-deleted when the instance is deleted. Valid values are: `false`, `no`, `true`, and `yes`.

For example, the following command creates an instance template based on `my-source-instance`, with the option to use the original image from `data-disk-a`, but set auto-delete to `true` and replace `data-disk-b` with a custom image.

```
gcloud compute instance-templates create my-instance-template \
  --source-instance my-source-instance \
  --configure-disk=device-name=data-disk-a,instantiate-from=source-image, \
  auto-delete=true
  --configure-disk=device-name=data-disk-b,instantiate-from=custom-image, \
  custom-image=projects/coreos-cloud/global/images/coreos-alpha-1492-3-0-v201708
```

The following table shows the options for overriding how disks are defined in the template.

Disk type	Options
-----------	---------

Disk type	Options
Boot disk	<ul style="list-style-type: none"> • [Default] Use the same source image or image family that was used to create the boot disk in the source instance. • Use the URL of any (custom or public) image.
Other read/write persistent disks	<ul style="list-style-type: none"> • [Default] Use the same source image/source image family that was used to create the disk in the source instance. Note: If the source instance's disk doesn't have a source image/source image family property, then it is included in the template as a blank disk. • Use the URL of any (custom or public) image. • Use a blank disk in the template instead. When the template is used to create a new instance, this disk is created unformatted. You must <u>format and mount</u> (/compute/docs/disks/add-persistent-disk#formatting) the disk in a startup script before you can use it in a scalable setup. • Do not include the disk.
Read-only disk(s)	<ul style="list-style-type: none"> • [Default] Include the disk in read-only mode. • Do not include the disk.
Local SSD(s)	<ul style="list-style-type: none"> • [Default] Include a blank local SSD. When the template is used to create a new instance, this disk is created unformatted. You must <u>format and mount</u> (/compute/docs/disks/local-ssd#formatandmount) the disk in a startup script before you can use it in a scalable setup.

For each disk, you can also override the `auto-delete` attribute to specify whether the disk should be deleted when its associated instance is deleted.

By default, if no override options are specified, the disk configuration in the template matches the source instance.

Creating an instance template based on an existing template

You cannot update an existing instance template. But, if an instance template goes out of date or if you need to make changes, you can create another one with similar properties. Use the console.

1. Go to the **Instance templates** page.

[Go to the Instance templates page](https://console.cloud.google.com/compute/instanceTemplates) (https://console.cloud.google.com/compute/instanceTemplates)

2. Click on the instance template that you want to copy and update.
3. Click **Create similar**.
4. Update the configuration in the new template.
5. Click **Create**.

Creating an instance template with a container image

You can specify a container image in an instance template. By default, Compute Engine also includes in the template a [Container-Optimized OS image](/container-optimized-os/docs) (/container-optimized-os/docs) with Docker installed. When you use the template to create a new instance, the container is launched automatically as the instance starts up.

[Consolegcloud](#) (#gcloud)

1. Go to the **Instance templates** page.

[Go to the Instance templates page](https://console.cloud.google.com/compute/instanceTemplates) (https://console.cloud.google.com/compute/instanceTemplates)

2. Click **Create instance template**.
3. In the **Container** section, select the **Deploy a container image to this VM instance** checkbox.
4. Specify the **Container image** to use.
 - For example, you can specify `gcr.io/cloud-marketplace/google/nginx1:1.12` to select an NGINX 1.12 container image from Google Cloud Marketplace.
 - If you use a container image from Docker Hub, always specify the full Docker image name. For example, specify the following image name to deploy an Apache container image: `docker.io/httpd:2.4`.
5. Optionally, click **Advanced container options**. For more information, see [Configuring options to run your Container](/compute/docs/containers/configuring-options-to-run-containers) (/compute/docs/containers/configuring-options-to-run-containers).
6. Click **Create**.

Creating an instance template that specifies a subnet

Instance template commands have `--subnet` and `--region` flags that place new instances into the subnet of your choice. The `--subnet` flag requires the `--region` flag.

```
d compute instance-templates create INSTANCE_TEMPLATE_NAME \  
-region REGION \  
-subnet SUBNET_NAME_OR_URL
```

Replace the following:

- ***INSTANCE_TEMPLATE_NAME***: The desired name of the instance template.
- ***REGION***: The region of the subnet.
- ***SUBNET_NAME_OR_URL***: Either the name of the subnet or its URL.

★ **Note:** If you use the subnet name, Google finds that subnet in the target region (provided that it exists there); however, if you use the subnet URL, the instance template can only be used to create instances in the specific region that is associated with the subnet URL.

The following example creates a template called `template-qa` that only creates instances in the `subnet-us-qa` subnet.

```
d compute instance-templates create template-qa \  
-region us-central1 \  
-subnet subnet-us-qa  
  
ed [https://compute.googleapis.com/compute/latest/projects/PROJECT_ID/global/instanc  
MACHINE_TYPE PREEMPTIBLE CREATION_TIMESTAMP  
ate-qa n1-standard-1 2015-12-23T20:34:00.791-07:00
```

Using this template to create instances for a [managed instance group](#) (`/compute/docs/instance-groups`) (with or without [autoscaling](#) (`/compute/docs/autoscaler`)) automatically creates the instance in the specified region and subnet. This lets you control the subnet of new instances created for load balancing.

Using custom or public images in your instance templates

Because managed instance groups are designed to add and remove instances frequently, it is useful to [create a custom image](/compute/docs/images/create-delete-deprecate-private-images) and specify it in the instance template. Prepare your image with the apps and settings that your instances need so you don't have to manually configure those items on individual instances in the MIG.

Alternatively, you can create an instance template that uses a [public image](/compute/docs/images) and a [startup script](/compute/docs/startupscript) to prepare the instance after it starts running. Custom images are more [deterministic](/compute/docs/instance-templates/deterministic-instance-templates) and start more quickly than instances with startup scripts. However, startup scripts are more flexible and let you update the apps and settings in your instances more easily.

If the image you want to use belongs to a different project, you can still use that image in your instance template if the owner of that project grants the MIG access to the images. For more information, see [Granting a MIG access to images](#)

[compute/docs/images/sharing-images-across-projects#granting_a_managed_instance_group_access_to_image](/compute/docs/images/sharing-images-across-projects#granting_a_managed_instance_group_access_to_image)

Updating an instance template

It is impossible to update an existing instance template or change an instance template after it has been created. If an instance template goes out of date, or you need to make changes to the configuration, create a new instance template.

What's next

- Read tips to create [Deterministic instance templates](/compute/docs/instance-templates/deterministic-instance-templates).
- Create a [VM instance from an instance template](/compute/docs/instances/create-vm-from-instance-template).
- Create a [zonal MIG](/compute/docs/instance-groups/creating-groups-of-managed-instances).

- Create a regional MIG
([/compute/docs/instance-groups/distributing-instances-with-regional-instance-groups](https://cloud.google.com/compute/docs/instance-groups/distributing-instances-with-regional-instance-groups)).

Except as otherwise noted, the content of this page is licensed under the [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/) (<https://creativecommons.org/licenses/by/4.0/>), and code samples are licensed under the [Apache 2.0 License](https://www.apache.org/licenses/LICENSE-2.0) (<https://www.apache.org/licenses/LICENSE-2.0>). For details, see the [Google Developers Site Policies](https://developers.google.com/site-policies) (<https://developers.google.com/site-policies>). Java is a registered trademark of Oracle and/or its affiliates.

Last updated 2020-07-30 UTC.