

[Security & Identity Products](https://cloud.google.com/products/security/) (<https://cloud.google.com/products/security/>)

[Resource Manager](https://cloud.google.com/resource-manager/) (<https://cloud.google.com/resource-manager/>)

[Documentation](https://cloud.google.com/resource-manager/docs/) (<https://cloud.google.com/resource-manager/docs/>) [Guides](#)

# Resource locations supported services

## Beta

This product or feature is in a pre-release state and might change or have limited support. For more information, see the [product launch stages](https://cloud.google.com/products/#product-launch-stages) (<https://cloud.google.com/products/#product-launch-stages>).

The resources used by each service are impacted by location in different ways. Before adding a [resource locations](https://cloud.google.com/resource-manager/docs/organization-policy/defining-locations)

(<https://cloud.google.com/resource-manager/docs/organization-policy/defining-locations>) constraint to your [organization policy](https://cloud.google.com/resource-manager/docs/organization-policy/overview)

(<https://cloud.google.com/resource-manager/docs/organization-policy/overview>), review the appropriate section below to see how the resources you are applying the policy to will behave.

**Note:** Organization policies, including those with the resource locations constraint, are not retroactive. Adding the resource locations constraint to existing resources will not restrict their location, but might restrict automatic child resource creation or other functionality.

## App Engine

in App Engine is a property of the `application` resource. The location property is enforced for all environments when you create an `application`. You can create only one App Engine `application` in each project. A Cloud Storage bucket is automatically created in the same location as the `application`. If you create an `application` with a broad location that doesn't comply with the organization policy, you will have to create a new project and App Engine `application`.

When you disable an `application`, it won't be served in future, but replicated code and data will remain in locations the `application` was stored in. To completely erase this data, delete the parent project.

The App Engine flexible environment is built on top of Compute Engine. Auto-scaling instances can fail if any locations where scaling happens aren't in the list of allowed locations that are defined in the organization policy.

For a list of available locations, see [App Engine Locations](https://cloud.google.com/appengine/docs/locations) (<https://cloud.google.com/appengine/docs/locations>).

## Compute Engine

Compute Engine uses a large number of resources, and these can be global, regional, or zonal. Regional and zonal resources are subject to the resource locations constraint. For example, you can use Compute Engine to create a virtual machine, and its physical location can be restricted by the resource locations constraint. Cloud Router, CDN Interconnect, and Cloud VPN resources are also regional and can be restricted by this constraint.

Some resources, such as Compute Engine snapshots and Compute Engine images, are global, but have the resource locations constraint enforced on them. A snapshot resource is created in a particular location decided by the user, or in an automatically selected location if one has not been specified. For more information about the location of Compute Engine resources, see [Global, Regional, and Zonal resources](https://cloud.google.com/compute/docs/regions-zones/global-regional-zonal-resources) (<https://cloud.google.com/compute/docs/regions-zones/global-regional-zonal-resources>).

Enforcement of the resource locations is handled at the level of the Compute Engine resource at the time you create it. The enforcement isn't retroactive: it only restricts the creation of resources in the listed regions, and doesn't restrict against movement of data.

Global resources, such as Compute Engine images and Cloud DNS resources, aren't subject to location-based enforcement. Note that some services, such as Virtual Private Cloud and Cloud Load Balancing, use both global and regional resources. For example, VPC networks are global, but subnetworks are regional and subject to resource locations constraint enforcement.

[VM instance groups](https://cloud.google.com/compute/docs/instance-groups/) (<https://cloud.google.com/compute/docs/instance-groups/>) are subject to location-based enforcement. Instance groups can be zonal (deployed to a single zone) or regional (deployed to multiple zones in the same region). Regional resources that enforce the resource locations constraint require that all zones provided by the user or selected by default match the defined locations. For example, if only one zone from a given region is restricted, then the user can still create a regional instance group, but must [select zones](#)

(<https://cloud.google.com/compute/docs/instance-groups/distributing-instances-with-regional-instance-groups#selectingzones>)

that are not location restricted.

When you use an instance group, the effective organization policy is evaluated on the creation of each instance. This will affect [autoscaling](https://cloud.google.com/compute/docs/autoscaler/), [autohealing](https://cloud.google.com/compute/docs/instance-groups/autohealing-instances-in-migs), [auto-updating](https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups)

(<https://cloud.google.com/compute/docs/instance-groups/rolling-out-updates-to-managed-instance-groups>)

, and manually [recreating](https://cloud.google.com/compute/docs/instance-groups/working-with-managed-instances#recreate_instances)

([https://cloud.google.com/compute/docs/instance-groups/working-with-managed-instances#recreate\\_instances](https://cloud.google.com/compute/docs/instance-groups/working-with-managed-instances#recreate_instances))

the instances or increasing the size of the instance group.

For a list of available locations, see the Compute Engine [Regions and Zones](https://cloud.google.com/compute/docs/regions-zones)

(<https://cloud.google.com/compute/docs/regions-zones>) page.

## Google Kubernetes Engine

Google Kubernetes Engine uses [Compute Engine regions and zones](https://cloud.google.com/compute/docs/regions-zones)

(<https://cloud.google.com/compute/docs/regions-zones>). Enforcement of the resource locations is handled at the level of the Compute Engine resource when you create the VM for a cluster. If you want to scale a cluster by adding more instances or adding another zone, these new additions must also be in an allowed location.

To create clusters with sufficient redundancy, use [value groups](#) (`#value_groups`) to control the locations that are restricted. If you set the locations manually, all zones in that region must be in the allowed list to have the same level of redundancy. Auto-scaling clusters can break if any of the locations in which scaling happens aren't in the list of allowed locations defined in the organization policy.

## Cloud Bigtable

A Cloud Bigtable instance resource is a logical container of clusters. Each of these clusters is located in a zone. All data in an instance is replicated uniformly to all clusters that are contained in that instance. Organization policy is enforced when a cluster is created. You can't

create new storage containers in a location that is denied in the organization policy. Existing instances and clusters will continue to operate even if they are in locations that are denied by the organization policy.

For a list of available locations, see the [Cloud Bigtable Locations](https://cloud.google.com/bigtable/docs/locations) (<https://cloud.google.com/bigtable/docs/locations>) page.

## Datastore

The Datastore **database** resources are directly dependent on the App Engine application in the parent project and its defined location. Disabling the App Engine application will block API access for the associated database. To delete replicated data from the physical locations, delete the project as described in the [App Engine](#) (#app\_engine) section.

For a list of available locations, see the [Datastore Locations](https://cloud.google.com/datastore/docs/locations) (<https://cloud.google.com/datastore/docs/locations>) page.

## Cloud Spanner

Organization policy is enforced when you create an instance. Instances are either regional or multi-regional resources. If an instance is blocked by the resource locations organization policy, the only way to bring the resource into compliance is by deleting the instance. Instances that are blocked by the resource locations organization policy will still allow for reads, writes, and the creation of database resources.

For a list of available locations, see the [Cloud Spanner Instances](https://cloud.google.com/spanner/docs/instances) (<https://cloud.google.com/spanner/docs/instances>) page.

## Cloud SQL

Organization policy is enforced when you create an instance. The instance is a regional resource that will create a zonal database, for which the resource location is not enforced. When you create read replicas or database clones, you locate the new resources in the same region as the original, so the resource locations organization policy isn't enforced.

For a list of available locations, see the Cloud SQL [Instance Locations](https://cloud.google.com/sql/docs/mysql/instance-locations) (https://cloud.google.com/sql/docs/mysql/instance-locations) page.

## Cloud Storage

Organization policy is enforced when you create a **bucket** resource. **Bucket** resources are regional or multi-regional. **Object** resources can be added to any existing **bucket** even if the **object** is in a location that is denied by the resource locations organization policy. To ensure that your resources are in compliance with the resource locations organization policy, create new **bucket** resources after the organization policy is applied, and then migrate data from old **bucket** resources to the new ones.

For a list of available locations, see the Cloud Storage [Bucket Locations](https://cloud.google.com/storage/docs/bucket-locations) (https://cloud.google.com/storage/docs/bucket-locations) page.

## Persistent Disk

The organization policy is enforced when you create a **disk** resource, which can then be attached to virtual machines:

- After you create a zonal **disk** resource, you can attach it to virtual machine instances in the same zone.
- After you create a regional **disk** resource, you can attach it to virtual machine instances in one of the two zones the **disk** resides in.

Organization policy compliance isn't enforced retroactively. To enforce a new resource locations organization policy on existing **disk** resources, you need to delete the **disk** resources and then create them again with the organization policy applied to the parent resource.

For a list of available locations, see the Compute Engine [Regions and Zones](https://cloud.google.com/compute/docs/regions-zones) (https://cloud.google.com/compute/docs/regions-zones) page.

## BigQuery

BigQuery `dataset` resources can be both regional and multi-regional. Organization policy compliance isn't enforced retroactively. To enforce a new resource locations constraint on an existing `dataset`, delete the `dataset` resource and create it again with the organization policy applied to the parent resource.

You can create `Database` resources within a `dataset` resource with a location that is denied by the resource locations organization policy. The location of the `dataset` resource does not dictate the location of the `database` resource. To enforce a new resource locations constraint on an existing `disk`, delete the `disk` resource and create it again with the organization policy applied to the parent resource.

For a list of available locations, see the BigQuery [Dataset Locations](https://cloud.google.com/bigquery/docs/dataset-locations) (<https://cloud.google.com/bigquery/docs/dataset-locations>) page.

## Dataflow

The organization policy is enforced when you create a `job`. A `job` is a regional resource that uses both Cloud Storage and Compute Engine. You can configure Compute Engine workers to execute in a zone outside the region of the `job` by specifying the `zone` parameter. In this case, the Dataflow control plane will execute in the specified region, while data processing workers will execute in the specified zone. If you don't specify the zone of workers, the workers will be created within the region in which the `job` is configured to run.

If you don't specify the zone of the `job`, the location of the workers will be in one of the zones within the region in which the `job` is configured to run. Dataflow will select the zone based on the available capacity in the zone. All zones within the region of the `job` should be set as allowed values in the resource locations organization policy.

Auto-scaling clusters can break if any of the locations in which scaling happens isn't in the list of allowed locations defined in the organization policy.

For a list of available locations, see the Dataflow [Regional endpoints](https://cloud.google.com/dataflow/docs/concepts/regional-endpoints) (<https://cloud.google.com/dataflow/docs/concepts/regional-endpoints>) page.

## Dataproc

When you create a `cluster`, the organization policy is enforced based on the region you specify in the creation request. The location of a `job` is bound by the location of the `cluster` that is its parent when the `submit` method is called.

For a list of available locations, see the Dataproc [Regional endpoints](https://cloud.google.com/dataproc/docs/concepts/regional-endpoints) (<https://cloud.google.com/dataproc/docs/concepts/regional-endpoints>) page.

## Pub/Sub

The resource locations organization policy affects the locations in which messages published to a `topic` can be persisted at rest. The organization policy is enforced when you publish messages to a `topic`. Note that a `topic` is still a global resource that is accessible from anywhere in the world to authorized clients.

Changes to the organization policy are not automatically applied to existing `topics`. If a new resource locations organization policy denies a location where messages published to a `topic` are already stored, those messages will not automatically be moved.

For more information, see the Pub/Sub [Restricting Pub/Sub resource locations](https://cloud.google.com/pubsub/docs/resource-location-restriction) (<https://cloud.google.com/pubsub/docs/resource-location-restriction>) page.

## Key Management Service

KMS resources can be created in regional, dual-regional, multi-regional, or global locations. The organization policy will be enforced at the time you create that resource.

For more information, see the KMS [locations](https://cloud.google.com/kms/docs/locations) (<https://cloud.google.com/kms/docs/locations>) page.

---

*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 our [Site Policies](https://developers.google.com/terms/site-policies) (<https://developers.google.com/terms/site-policies>). Java is a registered trademark of Oracle and/or its affiliates.*

*Last updated December 4, 2019.*