

Product or feature is in a pre-release state and might change or have limited support. For more information, see the [product launch stages](#) (/products/#product-launch-stages).

Google Cloud uses [regions, subdivided into zones](#) (/docs/geography-and-regions#regions_and_zones), to define the geographic location of physical computing resources.

You specify a location for storing your Cloud Healthcare API data when you create a dataset. After you create the dataset, the location cannot be changed. Data within the dataset is stored at rest in the chosen location.

The location is tied to the dataset's identity and is a permanent part of the dataset's [resource name](#) (/healthcare/docs/concepts/api-structure#resource_names). All data stores within the dataset are assigned to the same region as the dataset.

There are two types of locations:

- A *regional location* is a specific geographic place, such as Tokyo. For more information, see [Regional resources](#) (/docs/geography-and-regions#regional_resources) on the Geography and Regions page.
- A *multi-regional location* is a large geographic area, such as the United States, that contains at least two regional locations. For more information, see [Multi-regional resources](#) (/docs/geography-and-regions#multi-regional_resources) on the Geography and Regions page.

The Cloud Healthcare API supports a subset of the [full list of Google Cloud locations](#) (<https://cloud.google.com/about/locations>).

The Cloud Healthcare API is available in the following regions:

| Region name | Region description |
|---|-----------------------------------|
| North America | |
| us-central1 (https://cloud.google.com/about/locations/?region=americas#region) | Iowa, USA |
| us-west2 (https://cloud.google.com/about/locations/?region=americas#region) | Los Angeles, USA |
| northamerica-northeast1 (https://cloud.google.com/about/locations/?region=americas#region) | Montréal, Canada |
| Europe | |
| europe-west2 (https://cloud.google.com/about/locations/?region=europe#region) | London, UK |
| europe-west4 (https://cloud.google.com/about/locations/?region=europe#region) | Netherlands |
| Asia | |
| asia-southeast1 (https://cloud.google.com/about/locations/?region=asia-pacific#region) | Singapore |
| asia-northeast1 (https://cloud.google.com/about/locations/?region=asia-pacific#region) | Tokyo, Japan |
| Multi-region name | Multi-region description |
| us (https://cloud.google.com/about/locations/?region=multi-region#region) | Data centers in the United States |

When you choose a location for your data, you might want to consider factors such as:

- Regulatory requirements about where to store your data
- Latency

- Resiliency
- Cost
- Colocation with other Google Cloud services

For example, Google manages [multi-regional locations](#) (/docs/geography-and-regions#multi-regional_resources) to be redundant and distributed within and across regions. These services optimize availability, performance, and resource efficiency. As a result, these services require a trade-off on either latency or the consistency model.

Consider doing the following when choosing a location for your data:

- **Colocate your dataset and your external data source.**
- **Colocate your dataset with your Cloud Storage buckets when importing data.**
- **Colocate your dataset with your Cloud Storage buckets and BigQuery datasets when exporting data.**

You cannot change the location of a dataset after it is created. Also, you cannot move a dataset from one location to another. If you need to move data from one location to another, complete one of the following processes:

1. [Export](#) (/healthcare/docs/how-tos/fhir-import-export#exporting_fhir_resources) the data from your FHIR stores to a regional or multi-regional Cloud Storage bucket. When you export the data, the operation only exports the current version of each resource. The operation does not export version history; there is no bulk export operation for version history.

Note that there are [charges](#) (/healthcare/pricing#et_operations) for exporting FHIR data to Cloud Storage. You also incur charges for [storing the exported data](#) (/storage/pricing#storage-pricing) in Cloud Storage.

2. After you transfer the data to a Cloud Storage bucket, create a new dataset in the new location. Create any FHIR stores in the new dataset that you require for storing your data. Then, [import](#) (/healthcare/docs/how-tos/fhir-import-export#importing_fhir_resources) your data from Cloud Storage into the new FHIR stores.

Note that there are [charges](/healthcare/pricing#request_volume) (/healthcare/pricing#request_volume) for importing FHIR data into the Cloud Healthcare API.

1. [Export](/healthcare/docs/how-tos/dicom-import-export#exporting_dicom_instances) (/healthcare/docs/how-tos/dicom-import-export#exporting_dicom_instances) the data from your DICOM stores to a regional or multi-regional Cloud Storage bucket.

Note that there are [charges](/healthcare/pricing#etl_operations) (/healthcare/pricing#etl_operations) for exporting DICOM data to Cloud Storage. You also incur charges for [storing the exported data](/storage/pricing#storage-pricing) (/storage/pricing#storage-pricing) in Cloud Storage.

2. After you transfer the data to a Cloud Storage bucket, create a new dataset in the new location. Create any DICOM stores in the new dataset that you require for storing your data. Then, [import](/healthcare/docs/how-tos/dicom-import-export#importing_dicom_objects) (/healthcare/docs/how-tos/dicom-import-export#importing_dicom_objects) your data from Cloud Storage into the new DICOM stores.

Note that there are [charges](/healthcare/pricing#request_volume) (/healthcare/pricing#request_volume) for importing DICOM data into the Cloud Healthcare API.

Unlike FHIR and DICOM, there are no import and export operations available for HL7v2 data. Instead, we recommend that you [get the contents](/healthcare/docs/how-tos/hl7v2-messages#getting_hl7v2_message_contents) (/healthcare/docs/how-tos/hl7v2-messages#getting_hl7v2_message_contents) of your HL7v2 messages and write them to a new HL7v2 store in the new dataset location. To do so, complete the following steps:

1. Write the contents of the HL7v2 messages in your HL7v2 stores to a file.
2. Create a new dataset in your new location. Create any HL7v2 stores in the new dataset that you require for storing your data. Then, [create](/healthcare/docs/how-tos/hl7v2-messages#ingesting_and_creating_an_hl7v2_message) (/healthcare/docs/how-tos/hl7v2-messages#ingesting_and_creating_an_hl7v2_message) the messages one at a time from the saved file and store them in the new HL7v2 store.