




GCP Solution Packages

Solution architectures with Pricing estimates
for some popular use cases



A thick, vertical orange bar with rounded ends, positioned in the upper right quadrant of the page.

Part 1 Configuration

A thick, horizontal blue bar with rounded ends, positioned on the left side of the page.

This is an example of building a 2TB capacity Windows file server environment on Google Cloud.

Persistent Disk snapshots are used for backups, and a capacity of 3TB is assumed for the purpose of maintaining multiple generations of backups.

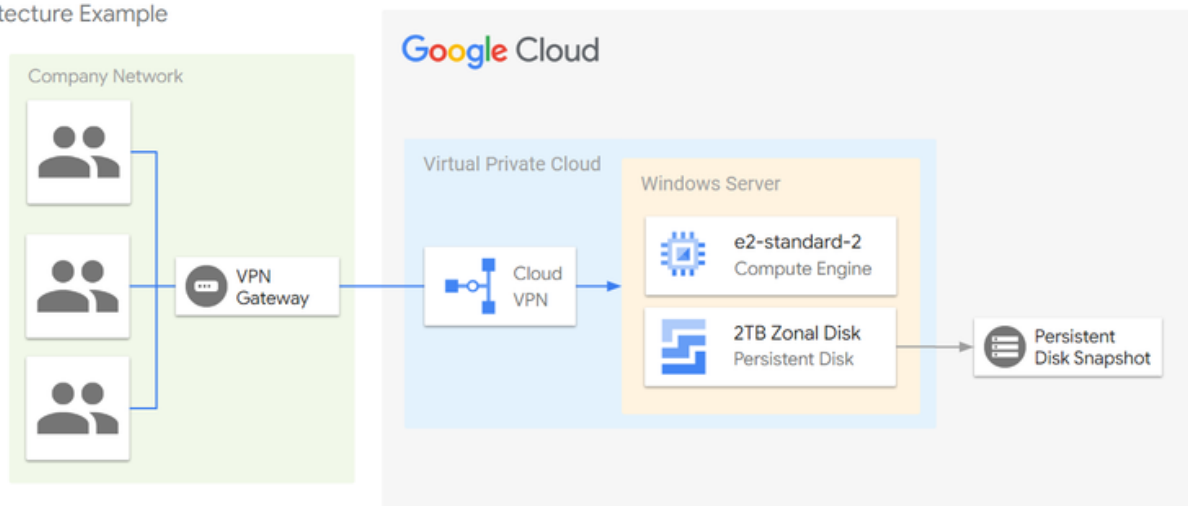
We assume that users will access the file server from the corporate network via VPN, and the transfer rate from Google Cloud to the corporate network is estimated to be about 10% of the data stored per month.

Active Directory, which provides user authentication, is not included in this example because there are multiple options, such as using the existing Active Directory on-premises or using Managed Microsoft AD on Google Cloud. For more information on this price estimate,

Google Cloud Solution Configuration

Windows File Server

Architecture Example



An example is a typical three-tier web application configuration where you can increase or decrease VM instances to meet your needs.

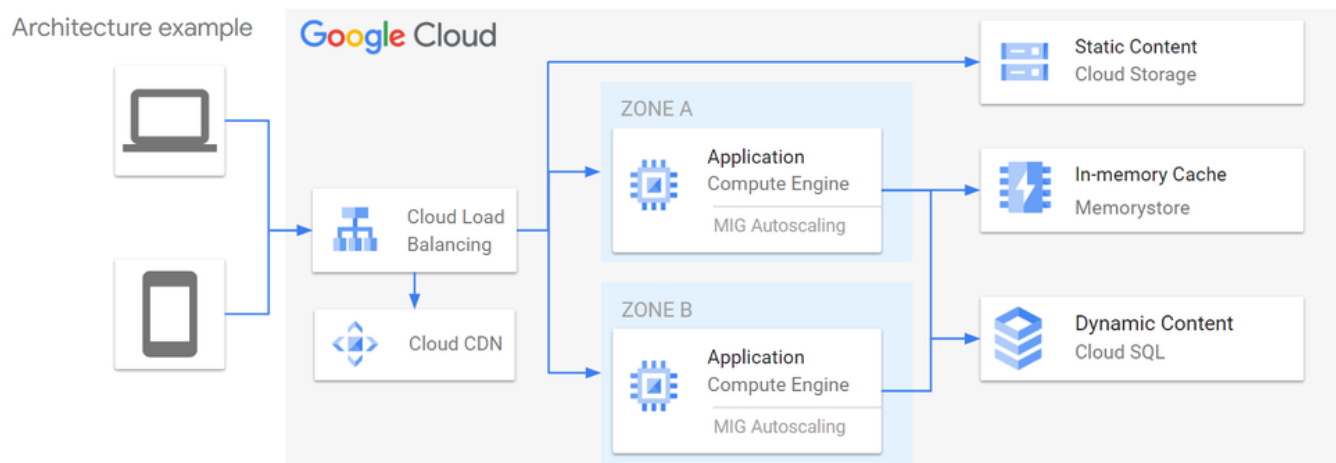
For automatic VM scaling, create a managed instance group (MIG) for the Compute Engine and then use the load balancer to take advantage of the global static external IP address to receive traffic.

In addition, we will reduce the load on the web server by utilizing caches such as Cloud CDN and Memorystore, and handling static files in Cloud Storage, which is an object storage.

This time, it is assumed that 3 instances of general e2-standard-2 (vCPU2, Memory 8GB) are running on the Compute Engine on average.

Google Cloud Solution Configuration

Data Analytics Platform



This is an example of analyzing data obtained from an on-premises environment or services other than Google with Google Cloud / Google Workspace.

Data to be analyzed will be sent to BigQuery from on-premises and external SaaS services via Google Cloud Storage once a day using BigQuery Data Transfer Service, with a total data capacity of 10 GiB per month, and data for the past 5 years shall be stored,

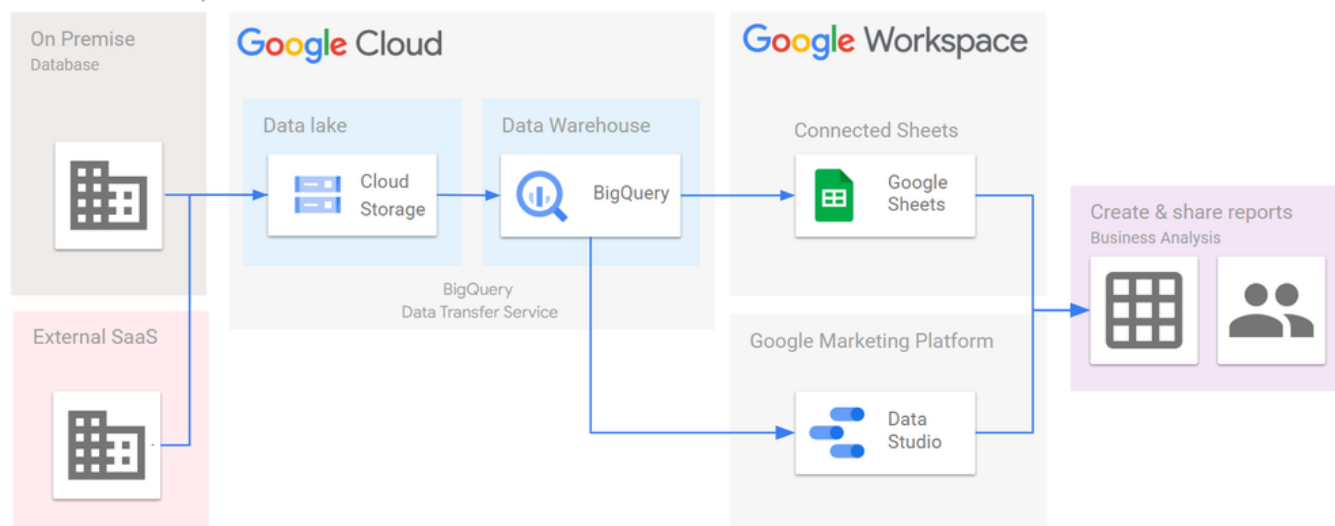
Users can visualize the data stored in BigQuery by creating reports using Google Sheets and Data Portal. The monthly analytic query executed by BigQuery assumes 500 GiB each month.

By storing data in BigQuery, you can easily integrate with the AI / ML services and advertising services provided by Google.

Google Cloud Solution Configuration

3-tier Auto Scaling Web Application Platform

Architecture Example



An example is a three-tier web application configuration where you can increase or decrease your Cloud Run instances to meet your needs.

Cloud Run autoscaling uses a global static external IP address to receive traffic using a load balancer with a minimum number of instances of 1 to eliminate cold start delays.

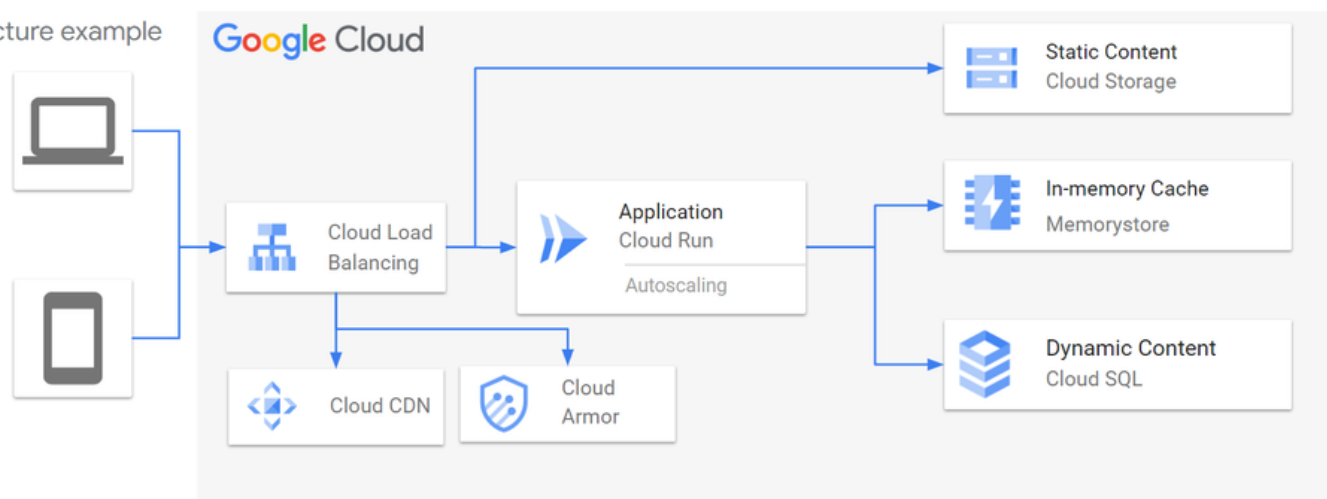
In addition, we will reduce the load on the web server by utilizing caches such as Cloud CDN and Memorystore, and handling static files in Cloud Storage, which is an object storage.

This time, we assume that the Cloud Run instance is 2 vCPU and Memory 512MB, and that 50 requests are processed simultaneously per instance.

Google Cloud Solution Configuration

3-tier Autoscaling Serverless Web Application Platform

Architecture example



For e-commerce sites, the challenges are to minimize response time and respond to the rapid increase in the number of accesses.

Static content such as HTML, JavaScript, CSS, and images are stored in Cloud Storage, and dynamic content is requested to compute services such as Compute Engine.

Compute Engine, which hosts dynamic content, is estimated to be processed by an average of two instances as a managed instance group (MIG).

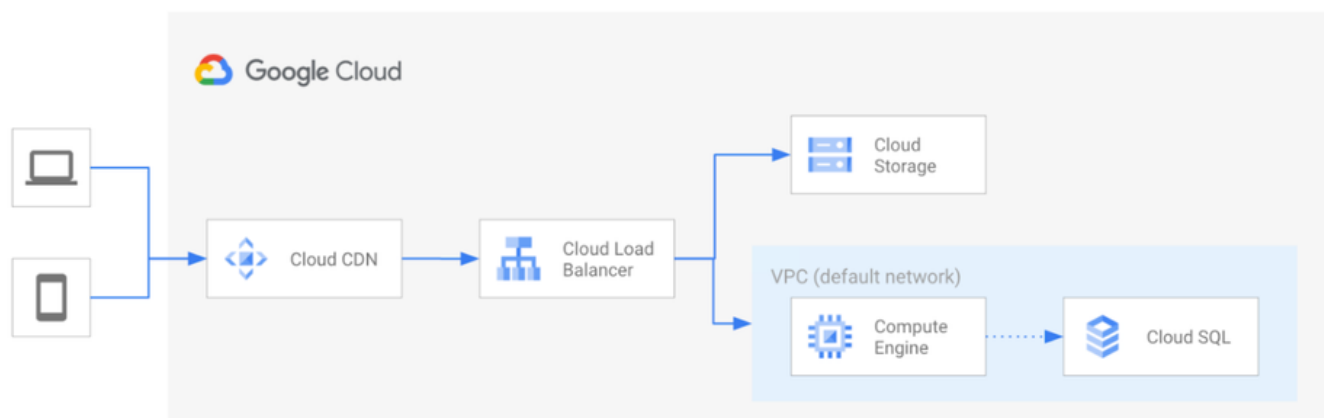
Cloud Load Balancing makes it easy to take advantage of cloud CDN caching of static content by simply enabling settings. In addition, Cloud Load Balancing's path routing function is used to determine whether to return the origin of static content or the origin of dynamic content depending on the path. This allows you to optimize cost and performance while minimizing requests to computing resources that generate dynamic content.

If you want more scalability in your database, consider Cloud Spanner as well.

Google Cloud Solution Configuration

EC Website Hosting

Architecture example



For e-commerce sites, the challenges are to minimize response time and respond to the rapid increase in the number of accesses.

Static content such as HTML, JavaScript, CSS, and images are stored in Cloud Storage, and dynamic content is requested to compute services such as Compute Engine.

Compute Engine, which hosts dynamic content, is estimated to be processed by an average of two instances as a managed instance group (MIG).

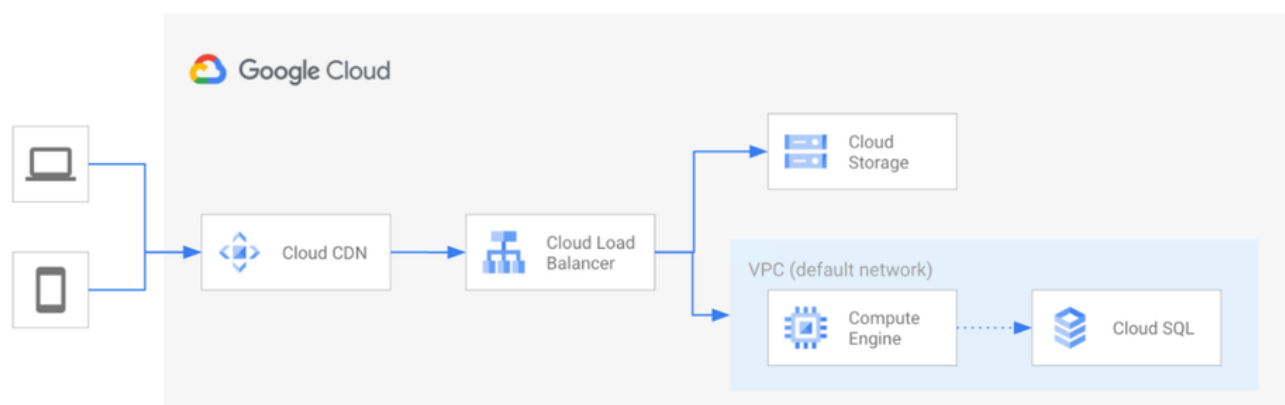
Cloud Load Balancing makes it easy to take advantage of cloud CDN caching of static content by simply enabling settings. In addition, Cloud Load Balancing's path routing function is used to determine whether to return the origin of static content or the origin of dynamic content depending on the path. This allows you to optimize cost and performance while minimizing requests to computing resources that generate dynamic content.

If you want more scalability in your database, consider Cloud Spanner as well.

Google Cloud Solution Configuration

EC Website Hosting

Architecture example



This is an example of a platform that analyzes various logs of applications running in other clouds and on-premises environments, as well as Google Cloud, centered on BigQuery.

This example assumes that you are running a containerized application in Cloud Run, while another component is running in another cloud or on-premises environment.

Logs for container applications on Cloud Run are aggregated to the default Cloud Logging. Set up Cloud Logging to automatically output log entries to BigQuery so that you can perform advanced analysis faster for a large number of aggregated logs.

Logs of applications running in other clouds or on-premises environments are also transferred and accumulated in Cloud Storage as a data lake, then ETL processed using Dataflow as needed, and saved in BigQuery.

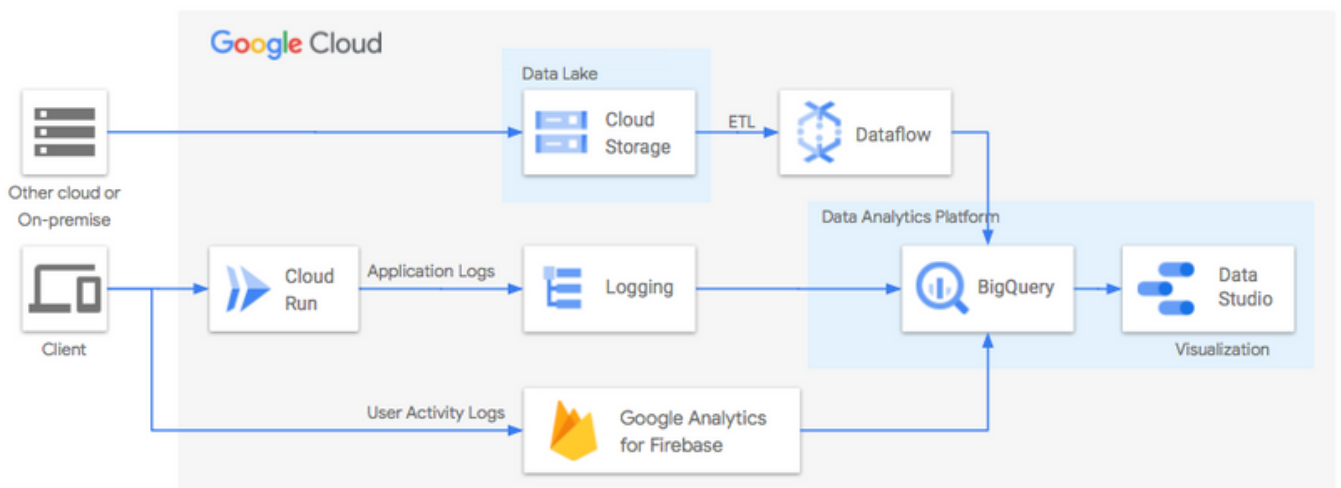
You can leverage Google Analytics for Firebase for user activity logs on your app. This log can also be integrated into BigQuery using standard features.

Various logs stored in BigQuery can be easily analyzed using SQL, and dashboards can be created using visualization tools such as Data Portal that can be easily linked.

Google Cloud Solution Configuration

Log Analysis Platform

Architecture example



This is an example of building a machine learning platform for batch prediction.

BigQuery ML allows you to build machine learning models and execute predictions using the SQL that data analysts are accustomed to, without having to move data from BigQuery, which normally stores data.

Using the model supported by BigQuery ML, it is possible to forecast demand, present recommended products, and classify users for campaign measures. Click [here](#) for model details.

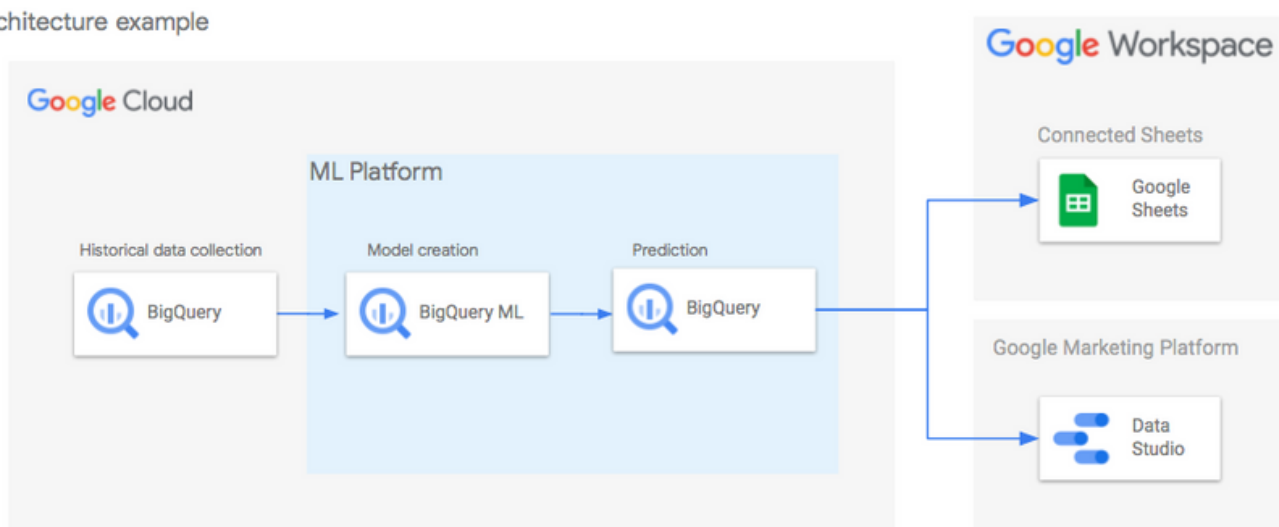
This is 5 stored in BigQuery

Based on this week's transaction information of 500 MiB from the data of the year, it shows an example that predicts the lifetime value of 1 million customers next week and leads to the consideration of campaign measures.

Google Cloud Solution Configuration

Log Analysis Platform

Architecture example



The Document AI API is a document understanding solution that takes unstructured data, such as documents and emails, and makes the data easier to understand, analyze, and consume.



This example is a simple architecture of a document processing pipeline that will automatically process documents that are uploaded to Cloud Storage. The pipeline consists of a primary Cloud Function that processes new files that are uploaded to Cloud Storage using a Document AI form processor and then saves form data detected in those files to BigQuery.

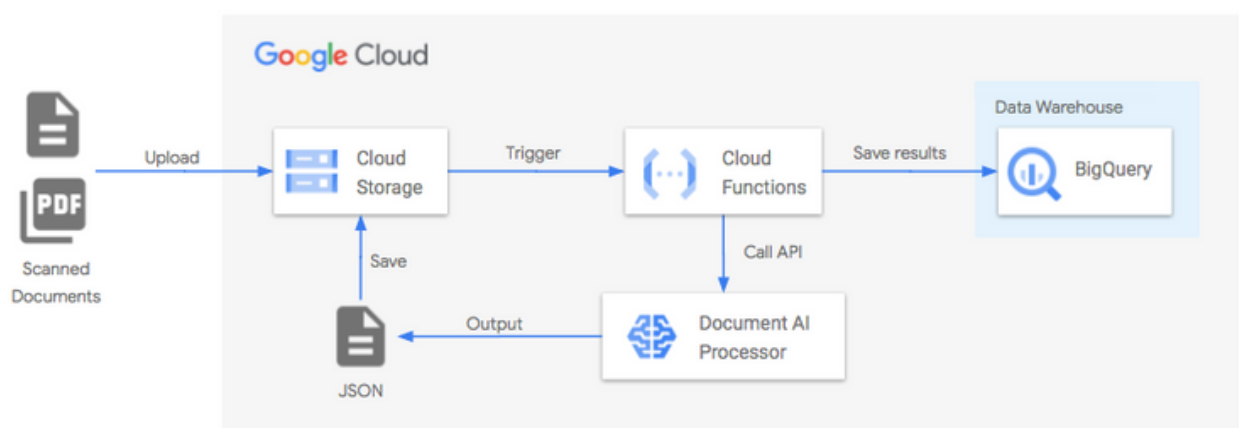
In this example the pipeline that uses a general form processor that will detect basic form data, such as a labelled field containing address information.

Document AI processors that use one of the specialized parsers that are beyond the scope of this example provide enhanced entity information for specific document types even when those documents do not include labelled fields. For example, a Document AI Invoice parser can provide detailed address and supplier information, from an unlabelled invoice document because it understands the layout of invoices.

Google Cloud Solution Configuration

Document Capture Pipeline using Document AI

Architecture example



Integrating Support API and Slack enable you to do following actions from Slack workspaces such as

- Post updates from Google Support to Slack
- Add comments to the case via Slack
- Find out how many cases are currently active under your organization

The pricing estimate is made on the following assumptions:

- Assuming 100 interactions / month with Support (case update)
- Cloud Run use “CPU is always allocated” mode

Steps and resources are available from

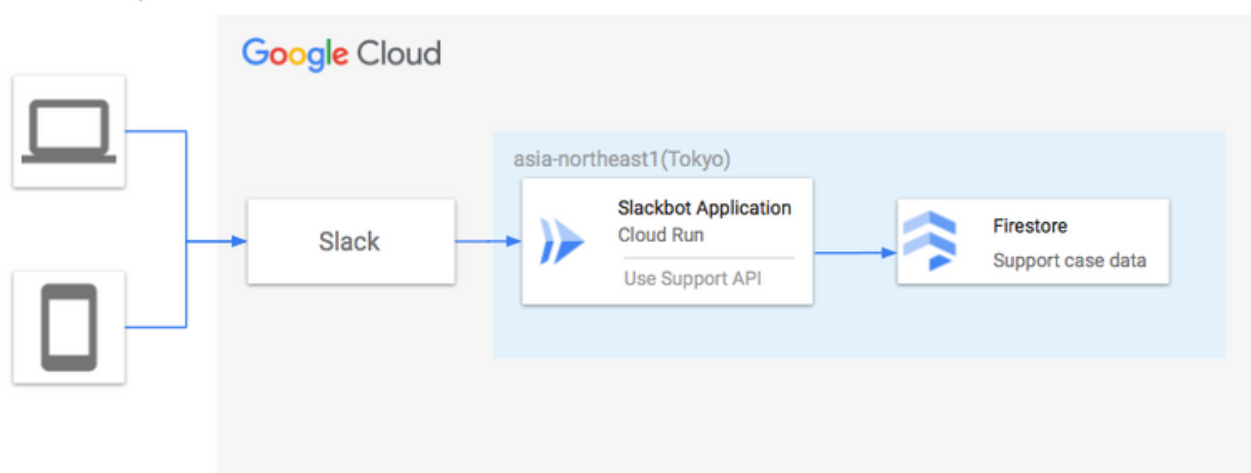
Slack Integration for Google Cloud Support in the Github repository:

<https://github.com/GoogleCloudPlatform/professional-services/tree/main/tools/google-cloud-support-slackbot>

Google Cloud Solution Configuration

Slack Integration for Google Support

Architecture example



- Landing zone design in Google Cloud
- Decide how to onboard identities to Google Cloud
- Decide the resource hierarchy for your Google Cloud landing zone
- Decide the network design
- Decide the security for your Google Cloud landing zone
- Recommend best practices ...

Google Cloud Solution Configuration

Landing Zone Design




Security controls



Ressource hierarchy

Google Cloud
Landing Zone



Network design



Identity provisioning

This is an example of visualising data obtained from an on-premises environment or services other than Google or any multi-cloud environments.

Data to be visualised will be sent to Looker from the source and can be used to create visualisation dashboards and utilise Lookers BI Engine to get predictive analytics and insights on the data.

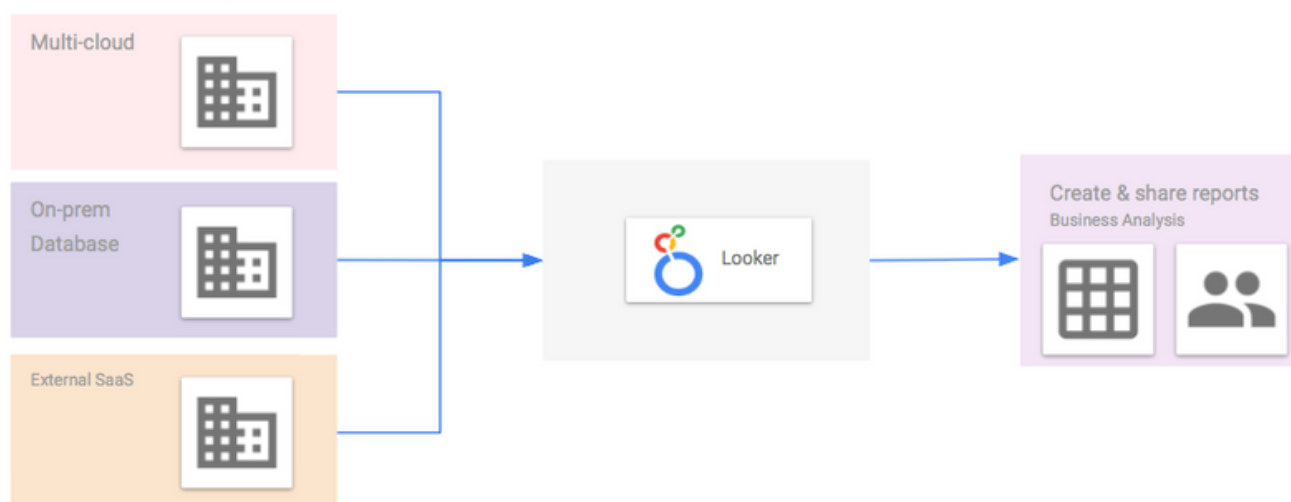
Users can also leverage Looker to build custom applications for data driven experience and design data driven workflows.

The Standard tier, which is the most popular choice among customers in the SMB space allows for up to 10k Query API access and 10k Admin API Access per month, Additional users and feature enhancing add-ons are available at an extra cost.

Google Cloud Solution Configuration

Data Visualization Platform

Architecture Example



This is an example of building an object detection pipeline for an internal online prediction application.

Vertex AI provides a unified platform for managing datasets, using prebuilt ML solutions, training custom models and deploying models for inference. With the AutoML Vision API, an object detection model can be trained with minimal knowledge of machine learning for images.

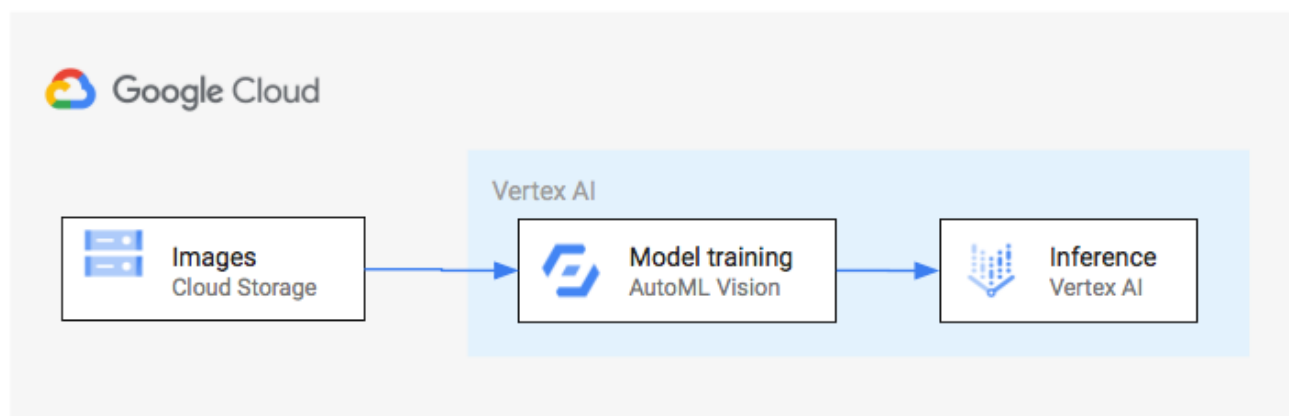
In this example, 500,000 images of 1MB each, stored in Cloud Storage, are used to train an AutoML Vision model for object detection (as an annual one-time task). This model is then deployed for online inference for 12 hours a day, 5 days a week.

Note that the costs for the one-time training are distributed over a 12 month period. Also, any charges for hosting and running an application that accesses the Vertex AI model for inferences is not considered here.

Google Cloud Solution Configuration

Online Object Detection

Architecture Example



This solution explores an example architecture pattern that ingests, processes and analyzes a large number of events concurrently from many different sources. The processing happens as events unfold, enabling you to respond and make decisions in real-time.

- Pub/Sub acts a pipeline for ingestion of real-time events.
- Dataflow performs the data transformation.
- That data is loaded into the BigQuery analytics engine.

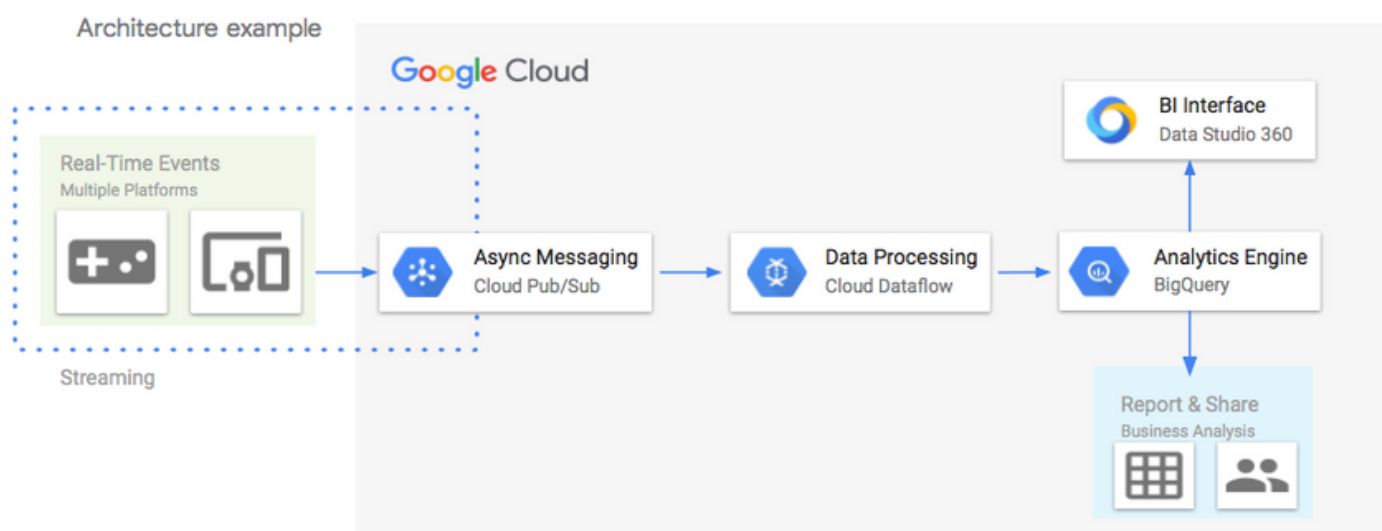
Data Studio or Looker can be used for creating dashboards and visualizing the data.

The pricing estimate is made on the following assumptions:

- 100 MB per user per hour
- 1000 users

Google Cloud Solution Configuration

Real Time Data Streaming and Analytics



This solution explores an example architecture pattern that deploys a new Wordpress website exposed to the public internet. This architecture can be used for the Blog, Intranet / internal Wiki and E-commerce platform.

Cloud Run: serverless PaaS offering to host containers for web-oriented applications, while offering security, scalability and easy versioning

Cloud SQL: Managed solution for SQL databases

VPC Serverless Connector: Solution to access the Cloud SQL VPC from Cloud Run, using only internal IP addresses

The necessary Terraform files are available from the following link to deploy a functioning Wordpress website with minimal technical overhead

Google Cloud Solution Configuration

Wordpress on Cloud Run

Architecture example

