



Google BigQuery

The Platform for Unlocking Real-time Business Insights

BigQuery is Google Cloud's fully managed, petabyte-scale, low-cost analytics data warehouse. BigQuery is NoOps, meaning there is no infrastructure to set up or manage and customers don't need a database administrator. This allows customers to focus on analyzing data to find meaningful insights, use familiar SQL, and transform their businesses with ease.

Business Challenge

Most large-scale organizations' current analytics systems do not adequately support digital transformation. Decision makers need the ability to quickly and cost effectively process and make use of their rapidly growing datasets.

Risk of Irrelevancy

Most large-scale organizations' current analytics platform does not adequately support timely decision making approach. Business leaders know they need to modernize the data ecosystem to serve their customers and gain competitive advantages.

Opportunity Cost of Unused Data

Data is hugely valuable, but organizations are unable to leverage it to drive business growth. Business leaders need a way to productize and monetize data quickly, without scores of engineers.

Inefficient and Data-Deprived Workforce

Employees at all levels waste time on manual tasks and make decisions - both large and small - with incomplete or inaccurate information.

BigQuery Overview and Key Capabilities

Google's mission statement is "to organize the world's information and make it universally accessible and useful". BigQuery meets this guiding principle by providing a cloud-native enterprise platform for data warehouse enabling secure access to near real-time data.

BigQuery is Google's fully managed, cost effective, serverless cloud data warehouse that scales with storage and computing power needs. With BigQuery, you get a columnar and ANSI SQL database that can analyze terabytes to petabytes of data at blazing-fast speeds. BigQuery's enterprise workload management can enable customers to get 100% predictability to manage complex workloads with self service experience.

Key capabilities include:

- **Serverless & Highly Scalable:** With serverless data warehousing, Google does all resource provisioning behind the scenes, so you can focus on data and analysis rather than worrying about upgrading, securing, or managing the infrastructure. With BigQuery's separated storage and compute, you have the option to independently scale the storage and processing tiers that make sense for your business and control access and costs for each. Customers in various

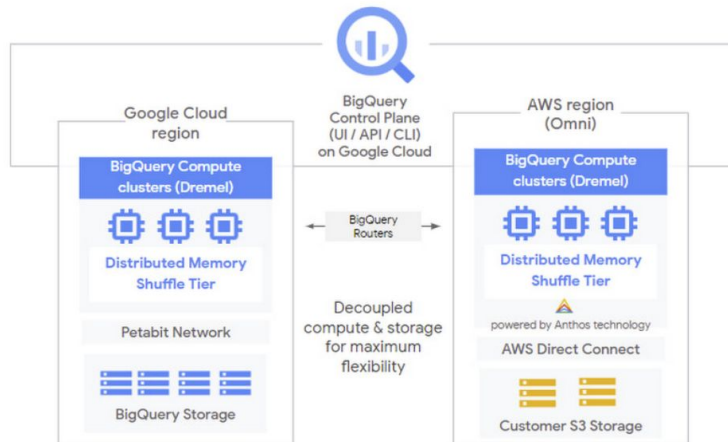
vertical industries leverage the scale of BigQuery to power all analytics. It can ingest millions of events into BigQuery and leverage it to make real-time decisions and for CAMELS scoring mechanism.

- **Integrated AI & ML:** BigQuery has embedded machine learning capability through BigQuery Machine Learning (BQML) to improve ratings and scoring of data. Integrated with best-in-class Cloud AI technologies like TensorFlow and TFX from Dataproc and Dataflow.
- **Highly Reliable Infrastructure:** A distributed analytic platform with separation of compute and storage which delivers unprecedented levels of reliability removing the burden of manageability from DBAs, compute and storage administrators.
- **Security & Compliance:** BigQuery provides strong security and governance controls with fine-grained controls through integration with Identity and Access Management. Rest assured knowing your data is encrypted at rest and in transit by default.
- **Ease of Use:** Google Analytics is known as the easiest platform to derive insights from data. The platform has integration with various BI Tools like Looker, Connected Sheets, and Excel. Provides Rich Semantic Modeling Layer for data modeling, which allows for layered architecture and seamless interoperation with 3P tools and expanded user reach for applications from ML to BI

Bringing multi-cloud analytics to your data with BigQuery Omni

A recent Gartner research survey on cloud adoption revealed that more than 80% of respondents using the public cloud were using more than one cloud service provider (CSP). While data is a critical component of decision making across organizations, for many, this data is scattered across multiple public clouds. BigQuery Omni is an extension of our continued innovation and commitment to multi-cloud, bringing you the best analytics and data warehouse technology, no matter where your data is stored.

BigQuery Omni is a flexible, multi-cloud analytics solution that lets you cost-effectively access and securely analyze data across Google Cloud, Amazon Web Services (AWS), and Azure (coming soon), without leaving the familiar BigQuery user interface (UI). Using standard SQL and the same BigQuery APIs our customers love, you will be able to break down data silos and gain critical business insights from a single pane of glass. And because BigQuery Omni is powered by **Anthos**, you will be able to query data without having to manage the underlying infrastructure.



Customer Case Studies



UPS uses BigQuery to achieve precise package volume forecasting for the company.



Twitter modernizes ad engagement analytics platform and aggregates millions of metrics per second.



The Home Depot nails its customer service with timely data that helps keep 50K items stocked across 2K stores.



Dow Jones unlocks business value of 1.3 billion news articles and uncovers hidden relationships and insights.

Advantages of BigQuery Platform

- ❖ **Increase operational efficiency**
 - Centrally manage diverse data sets at scale (Petabytes) and higher degree of performance
 - Establish a single source of truth to improve data accuracy
 - Highly available and reliable with hands-off setup
 - Well suited for both real-time and batch pipelines
 - Supports real-time data operations by native integration with streaming and API based data sources
- ❖ **Built-in AI / ML Capabilities**
 - Built-in ML models eliminates the need for data modelling
 - Support for user created ML models
- ❖ **Reduced costs**
 - Very little administration overhead
 - Reduce duplication of data between risk, regulatory, compliance, and customer intelligence systems

❖ **Increase value**

- Customize products based on enhanced knowledge of data
- Better understanding of data patterns, preferences, and behaviors
- Enable a more complete 360 degree view to gain data-driven capabilities and deeper analytics
- Integrated BI and SQL engine for ease of reporting

❖ **Integrated Data Security**

- Data encrypted at rest and in transit
- Immutable audit logs
- DLP API for redaction

Google Cloud Professional Services Offerings

Google Cloud's Professional Services Organization (PSO) is able to work directly with customers to help design, build, and deploy Google BigQuery to unlock the value of the data and cost-effectively process data analysis. The below outlines a services engagement to allow customers to assess BigQuery related offerings comprehensively.

Key activities	Deliverables	Engagement Details
<p>GCP data architecture deep-dive workshop</p> <ul style="list-style-type: none"> Review data architecture on Google Cloud concepts, options, considerations, and objectives to enable your team. <p>EDW workloads inventory and assessment</p> <ul style="list-style-type: none"> Conduct high-level analysis and document pertinent details of your existing EDW workloads, associated business and technical use cases, and key functional and nonfunctional requirements. Perform a ranking exercise on known workloads against cost and complexity. <p>Critical data sources inventory</p> <ul style="list-style-type: none"> Review all existing current- and future-state data integration points and assess criticality for up to three initial workloads targeted for the first iteration of your migration. <p>Architecture and design discussion</p> <ul style="list-style-type: none"> Facilitate a discussion focused on sharing best practices and arriving at key decision points on how to design and deploy your EDW target(s) onto Google Cloud. Some of the design topics include data storage, data modeling and schema design, ETL/ELT pipelines, job orchestration, logging and monitoring, dataflow job creation and templating, and error handling. <p>Recommendations</p> <ul style="list-style-type: none"> Capture and share recommendations around common obstacles or challenges to avoid, and other key considerations for determining an initial workload and subsequent workloads targeted for migration. 	<ul style="list-style-type: none"> Project charter - the goals agreed on with key stakeholders, and the expected outcomes of the migration program Deployment progress reports - reports on migration status at the beginning, middle, and end of the engagement Technical assessment report - aggregation of data and subsequent analysis of your existing EDW information into an actionable artifact to drive modernization efforts Reference architecture - reference architecture and high-level design documents depicting the solution recommended by Google's Professional Services to support the iterative, scalable migration of existing enterprise data warehouse Initial migration plan - a plan depicting next steps and a recommended approach for delivery of initial iteration of the migration 	<ul style="list-style-type: none"> Prerequisites: Cloud Foundations The timeframe for the engagement is dependent on project scope and agreed upon in the project charter The typical team consists of cloud consultants and subject matter experts with technical expertise based on your specific needs Work can be performed onsite or offsite at your discretion Your team will provide business decision makers, technical subject matter experts, and data stewards with knowledge of your existing EDW implementation(s) and access to the GCP resources required throughout the engagement Cost - Depending on the services, the solution costs \$100k-\$150k

Let's connect to discuss how Google's Analytics Platform can help your organization!

For more information visit <https://cloud.google.com/solutions/government>