BigQuery is getting Less Expensive for Data Driven Organisations that take action

Written by:

making science

Contributors:

Julio José Prado Senior Google Certified Data Scientist & Prof. Data Eng.

Kevin Daly Business Unit Director DataOps & Data Science

I. Background to the Challenge

Starting on July 5, 2023, BigQuery customers will be faced with an increase in the price of the on-demand "pay-per-byte" analysis model by 25% across all regions. Additionally BigQuery customers will no longer be able to purchase flat-rate annual, flat-rate monthly, and flex slot commitments.

Over the years we have seen the huge leaps made by Google Engineers to combine big data analytics and machine learning capabilities into an easy to use platform called BigQuery. BigQuery has always been a perfect fit for digital business analytics based on its proximity to the online world of consumers and the plug-and-play availability of Google analytics data and marketing services. The serverless functionality and performance gets better but customers are demanding more predictability with regard to performance and costs.

As a reseller of BigQuery and Google Cloud Platform for more than 7 years, we have experienced all types of data analytics and machine learning implementations using Google BigQuery. BigQuery is by far and away the data analytics cloud service that is easiest to use and delivers exceptional capabilities at a very competitive price. Sometimes larger enterprises find that with large amounts of data in BigQuery, large user bases and important processing demands, the cloud costs can start to add up quickly.

Google has recently announced a "capacity-commercial" model called BigQuery Editions. **BigQuery Editions** is going to make the Google Data warehouse and Data Analytics serverless platform **Faster, Better and Cheaper** securing its position as the best in class data analytics platform. But customers must take action with their data.

The customers will be required to perform capacity planning on their projected BigQuery utilisation. With the availability of more and more serverless IaaS in the cloud we all began to feel that an IT discipline known as Capacity Planning was now an unneeded skill-set, Google has just changed that premise for efficient use of serverless services.

Capacity Planning involves determining the appropriate amount of computing resources needed to support the expected workload and user activity, to ensure that the system can sustainably handle the data processing and analytics tasks efficiently and effectively.

In the case of Big Query, capacity planning is particularly

important because of the large volumes of data that are typically stored and processed within the system. The amount of computing resources required can vary greatly depending on factors such: as the volume and variety of data being loaded and warehoused, the volume and variety of data being regularly processed, complexity of the queries being run, the organisation of the data being queried and the number of ad-hoc users. The ad-hoc users is a particularly challenging workload to plan for, as the profiles and expertise level can vary greatly between Analysts and Data Scientists.

The diagram below from Google's VP and GM Data & Analytics, Gerrit Kazmaier, beautifully sums up the benefits of capacity planning with BigQuery Editions.

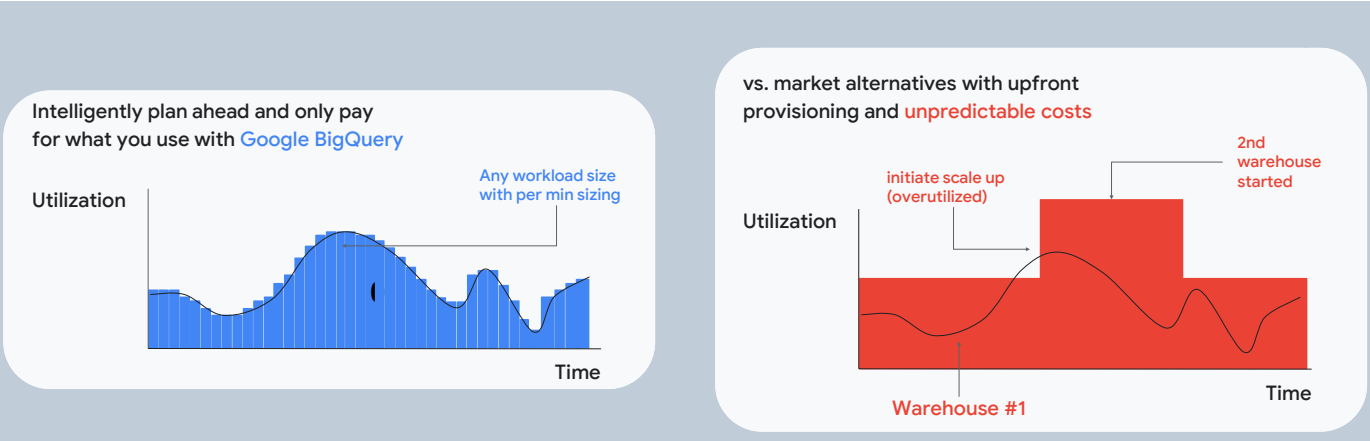

Google's **BigQuery On-Demand processing and pricing will still be available** to customers that do not want or have the capability to perform some capacity planning. But come the 5th of July 2023 BigQuery will be significantly more expensive for those customers that have not performed and configured their Google BigQuery projects with the necessary BigQuery Editions configurations.

II. Introduction to BigQuery Capacity-Commercial Models

BigQuery has two principal domains for capacity planning and commercial implications for each domain. The 1st domain is the Compute processing which is anything that processes the data in BigQuery and the 2nd domain is the Storage for all the data that is warehoused. The Compute processing domain is measured in "slots"and can be considered the equivalent of CPU+Memory to process a query, job, piece of a query etc. Generally speaking the more slots the faster your queries will run but keeping in mind that slots cost $$$ and it is ok if stages of a query/job are queued until slots become available; so you will be looking to right-size slots assigned to the given job.

The storage domain typically makes up 20% to 30% of the total monthly spend on BigQuery.

The BigQuery Storage commercial model is applied independently of the Compute processing model(s) selected.

IMPORTANT OPPORTUNITY: There is a new storage commercial option available worth analysing in your capacity planning activities whether your dataset commercial option should be Logical or Physical. Logical: This is the default billing model of a dataset. The data size is calculated based on the data types of the individual columns. For example, INT64 type takes 8 logical bytes. Physical: The data size is calculated based on the data stored on the disk after compression.

Active Logical Storage $0.02 per GB	**Long-term Logical Storage** $0.01 per GB	**Active Physical Storage** $0.04 per GB	**Long-term Physical Storage** $0.02 per GB
• Calculated on uncompressed data size	• Calculated on uncompressed data size • Automatic commercial classification change for data not modified in 90 days	• Calculated on compressed data size	• Calculated on compressed data size • Automatic commercial classification change for data not modified in 90 days

The BigQuery On-demand Compute processing model will continue to be available just more expensive than in the past and without the flex-slots availability.

On-Demand $5.00 per TB processed
Most expensive option for SQL queries and Analytics - Soft capped at 2000 concurrent slots - A maximum bytes billed setting to limit query costs - Query Editor available to estimate the number of bytes to process

BigQuery provides three new Compute Processing models known as BigQuery Editions (Standard, Enterprise, and Enterprise Plus) to support the compute processing workloads. These Editions provide an alternative to existing reservation models, such as flat-rate(soon to be deprecated) or on-demand models. BigQuery editions allow you to pick the right feature set for individual workload requirements. For example, the Standard Edition is best for development, and test workloads, while Enterprise has increased security, governance, machine learning and data management features. Enterprise Plus is targeted at very large enterprises that have requirements regarding encryption keys management and guaranteed service levels for given workloads or other specific regulatory requirements as noted. The table below describes each BigQuery Edition option as published on 29th of March 2023.

Standard

$0.04/slot hour

Low-cost option for standard SQL analysis

- Capped at 1,600 slots per reservation
- Max 5 reservations per admin project
- Autoscaling
- 99.9% SLA
- Zonal High Availability
- Google certifications[1] (Foundational compliance)
- Covered by Google Cloud BAA (for HIPAA compliance)
- Google managed keys

[1] All Google certifications including ISO 9001, ISO 27001, SOC 1-3, PCI

Enterprise

$0.06/slot hour

Advanced enterprise workloads

- **Standard +**
- Unlimited reservation size
- 99.99% SLA
- BI query acceleration
- Integrated ML models
- Full-text search
- Object tables
- VPC Service Controls to prevent data exfiltration
- Data masking, column security and row security

Optional 1 yr commitment ($0.048/slot hour)
Optional 3 yr commitment ($0.036/slot hour)

Enterprise Plus

$0.10/slot hour

Business critical enterprise workloads

- **Enterprise +**
- Regional-level disaster recovery[2]
- Customer-managed encryption keys
- Support for FedRAMP, ITAR and other compliance regimes available through Assured Workloads

Optional 1 yr commitment ($0.08/slot hour)
Optional 3 yr commitment ($0.06/slot hour)

[2] Roadmap functionality

On-demand (pay-as-you-go for data processed)

$6.25/TB scanned. First 1 TB/month is free

Includes all capabilities of Enterprise Plus

Google Cloud

You can choose to use BigQuery Editions and BigQuery On-demand models at the same time for your GCP Organisation. The challenge is deciding and planning on which Edition and/or other environment and pricing model when.

III. Benefits of BigQuery Capacity-Commercial Models

The first capacity commercial model to explore is the **Logical or Physical Storage** options for your BigQuery dataset. Google is giving most customers a sizable opportunity to save $$$ in their BigQuery storage costs.

The amount of disk storage space and subsequent cost that can be saved by selecting the physical storage option based on current pricing only makes sense if your data will compress by more than 50% as Google has doubled the storage cost for compressed data sets in BigQuery.

The level of compression of data depends on two main factors, the type of data and the compression algorithm used; sorry, you do not have control over the compression algorithm to be used in BigQuery at this writing.

In general, compression can result in significant BigQuery space savings, especially for large datasets. It is not uncommon for compression to reduce the size of a file or dataset by more than 50%.

For example, compressing text files (e.g ecommerce events, sales transactions, contact events, call logs, log files, documents, emails, etc) can often result in very significant BigQuery space savings closer 70% since text files tend to have a lot of redundancy and patterns that can be compressed.

On the other hand, compressing already compressed files, such as JPEG images or MP3 audio files, may not result in much additional space savings, since these formats already use compression techniques to reduce their size.

The examples below illustrate the benefits for BigQuery datasets supporting customers in different industries/sectors:

	Retail Sector Client			Education Sector Client			Travel Sector Client		
	Active Logical in (TB's)	Avg. Compression Result (in %)	Active Physical in (TB's)	Active Logical in (TB's)	Avg. Compression Result (in %)	Active Physical in (TB's)	Active Logical in (TB's)	Avg. Compression Result (in %)	Active Physical in (TB's)
Monthly Costs for BQ Storage for DWH & CDA domains	1250	72%	350	585	67%	193.05	985	73%	265.95
	$25,000.00		$14,000.00	$11,700.00		$7,722.00	$19,700.00		$10,638.00

IV. Specific Challenges and some solutions for defining a BigQuery capacity plan

With our years of experience in conducting capacity planning it is never as easy as you initially think; some examples of the challenges are shown below:

- Customer's BigQuery processing patterns occur over extended periods of time
- Cloud Monitoring will give us visibility on only 30 days of usage data unless customer is actively storing information
- BigQuery Information Schema statistics data is retained for: query statistics for 180 days, table statistics for 90 days.
- Google's BigQuery Estimator tool in preview limited to 30 days for On-demand configurations and only 7 for Editions configurations
- The number of workload domains to summarise usage and subsequent capacity planning
- Customer's have complex intercompany and intra-company billing / chargeback requirements.
- Customer's are forced to forecast for new lines of business, new user communities, new workloads as there is no historical data

Of course for each challenge there are methods to allow the capacity planners to push forward with a reasonable capacity forecast. At Making Science we have a methodical approach to capacity planning as introduced in the diagram below.

Making Science BigQuery Capacity Planning Methodology

- Data Collection is a critical part of BigQuery capacity planning. The most reliable scenario is basing the capacity plan based on historical data.
- In many cases the historical data does not offer a complete picture of the required capacity and data assumptions must be built by hand with subject matter experts in the business.

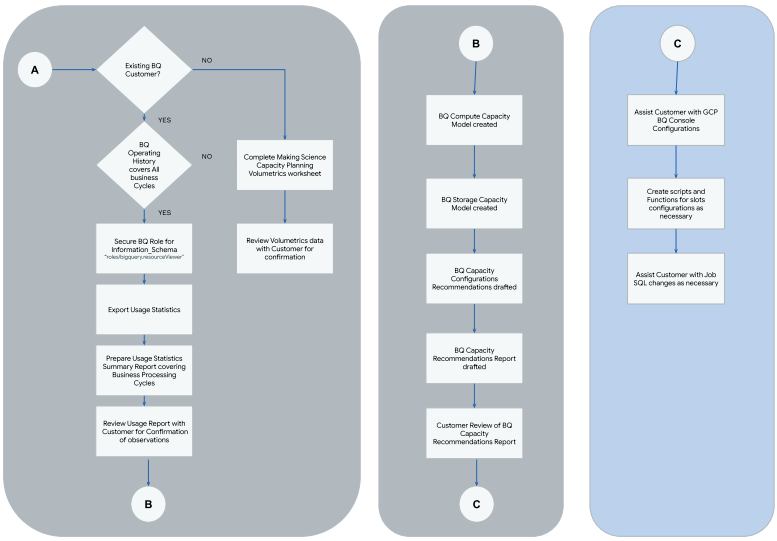

The flows help our customers and stakeholders understand the steps they need to take in conducting a thorough capacity planning activity for BigQuery. For most of our large enterprise customers, the capacity planning projects involve multiple customer departments, the Making Science DataOps team and very typically the Google Technical Account Manager. The RACI matrix helps all parties understand roles and responsibilities in the quick moving project.

Making Science BigQuery Capacity Planning Methodology

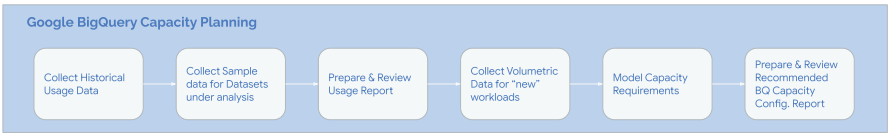

Activity	Customer	Making Science	Google
Set-up IAM access & BQ Role Resource.Viewer	R, A	C, I	-
Complete scoping Questionnaire	R, A	C, I	I
Collect BQ Historical usage data	C, I	R, A	-
Collect representative sample data & make available to MS	R, A	C, I	-
Prepare Usage Report	I	R, A	-
Walk-thru Usage Report & Volumetric data	C	R, A	I
Provide Intra-Company, intercompany billing requirements	R, A	C	I
Provide Volumetric data for any Missing and New Workloads	R, A	C	I
Model BQ Capacity Requirements (Compute & Storage)	I	R, A	C
Prepare BQ Capacity Planning Recommendations Report	-	R, A	C
Walk-thru BQ Capacity Recommendations Report	I	R, A	I

R = Responsible, A = Accountable, C = Consulted , I = Informed

We are firm believers in time-boxing the capacity planning project to ensure that the resulting configurations can be rolled out quickly and monitored thereafter to evaluate the GCP billing impact and make any additional adjustments as required. The diagram below offers draft planning for BigQuery capacity planning for very large enterprise customers.

Making Science BigQuery Capacity Planning timeline

This timeline is meant for orientation purposes only. Once the Scope of the BigQuery capacity planning exercise is defined with the inventory of GCP Organizations, GCP BQ projects and BQ Workloads a customer specific SOW and timeline can be closed.

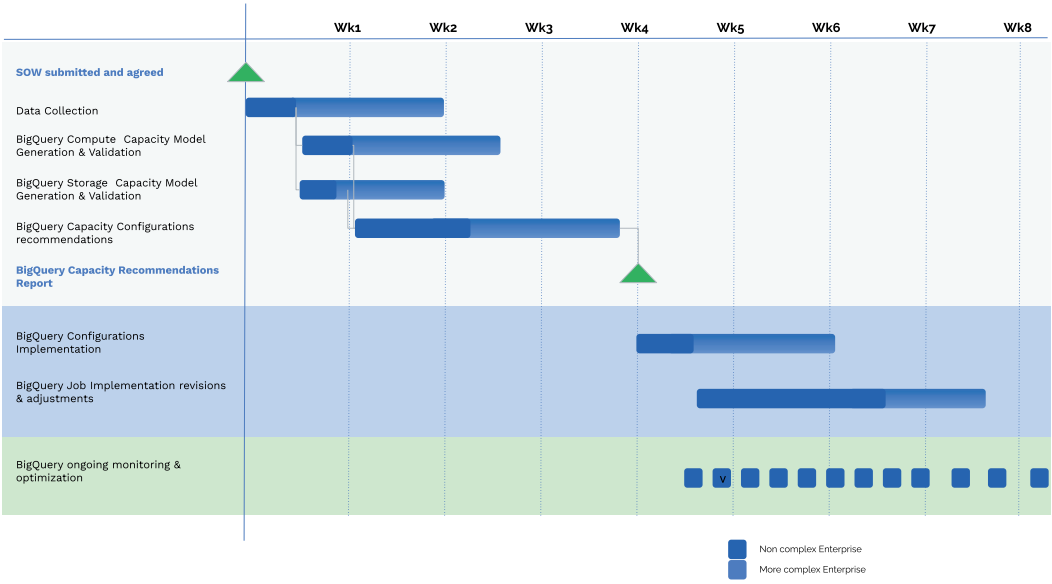

Smaller Enterprise customer capacity planning projects can typically be managed within a couple of weeks of time.

BigQuery Storage Capacity Planning

BigQuery offers an incredibly robust and resilient storage system as depicted in their architectural diagram below and now BigQuery customers have an opportunity to save $$$ on BigQuery storage costs:

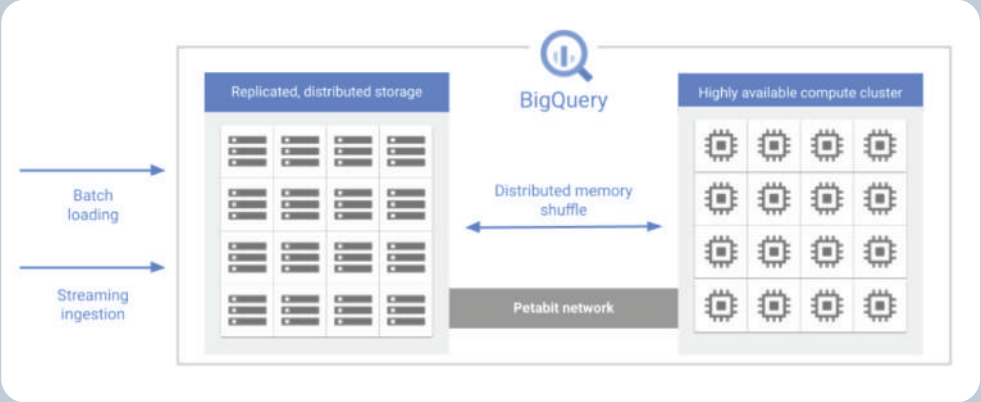

In general, the best way to determine whether we want to select a BigQuery Logical or Physical storage pricing option is to perform BigQuery store tests with two datasets. One dataset is configured with logical pricing and the other with physical pricing and we load data for the representative data sources and measure the compression benefits on a sample data and compare the size of the "physical table to the size of the "logical" table.

At Making Science our Engineers complete a dataset capacity planner. It is easiest to start with the existing GCP projects, datasets and tables.

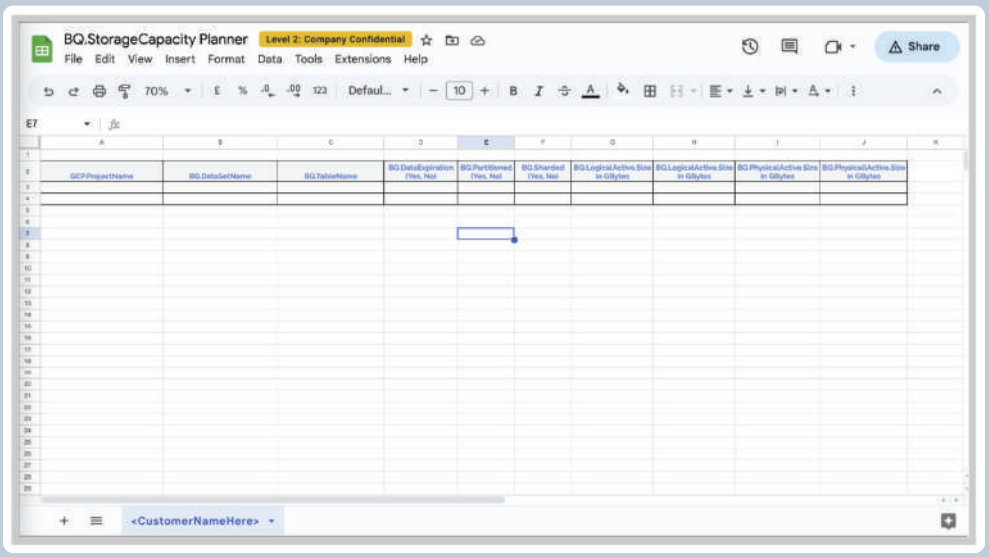

The engineers take a look at the existing configuration and logical definition and storage consumption by utilising the meta data in the **Information.Schema regarding tables.** The existing datasets (samples) are copied to new temporary datasets defined with the physical pricing option to compare the results.

This will give the guidance for the most cost effective BigQuery storage option to choose for the given dataset. We should keep in mind the default pricing option is "Logical" and that the pricing configuration is selected at the BigQuery Dataset level not the BigQuery table. At the time of this writing the only means to change the storage pricing option was by using an BigQuery alter command:s

```
ALTER SCHEMA DATASET_NAME
```

```
SET OPTIONS (storage_billing_model = 'physical');
```

CAUTION: This change can NOT be undone.

When you change a dataset's billing model, it takes 24 hours for the change to take effect.

In addition to selecting the best storage pricing option for each dataset, Making Science highly recommends a practice of tables organisation to ensure that there is the largest % of historical data being classified as "long-term" storage to receive very significant reductions in BigQuery storage costs. This may also require an audit/review of the SQL being executed to ensure SQL is properly written to receive standard SQL optimizations and the reductions in BigQuery costs. This is extremely important with on-demand pricing but also will have a significant impact when using BigQuery Editions.

BigQuery Editions Compute Capacity Planning

In BigQuery Editions, the unit used to measure the workload of a BigQuery job is the slot. A **slot** represents a virtual CPU used by BigQuery. VM slots are essentially units of computing power and memory that are used to run queries. Generally, if you utilise more slots, you can run more concurrent queries, and complex queries can run faster. BigQuery slot **Reservations** are the maximum pool size for Active slots usage to support concurrently running workloads (queries, pipelines, jobs, ML). The reserved slots can then be allocated down to folders, projects, and workloads as **Assignments** in effect defining a slots quota for the workload.

IMPORTANT NOTE: If a BigQuery job is started from a project that is assigned to a reservation, the job uses that reservation's slots. If a project is not assigned to a reservation (either directly or by inheriting from its parent folder or organisation), the jobs in that project use on-demand pricing.

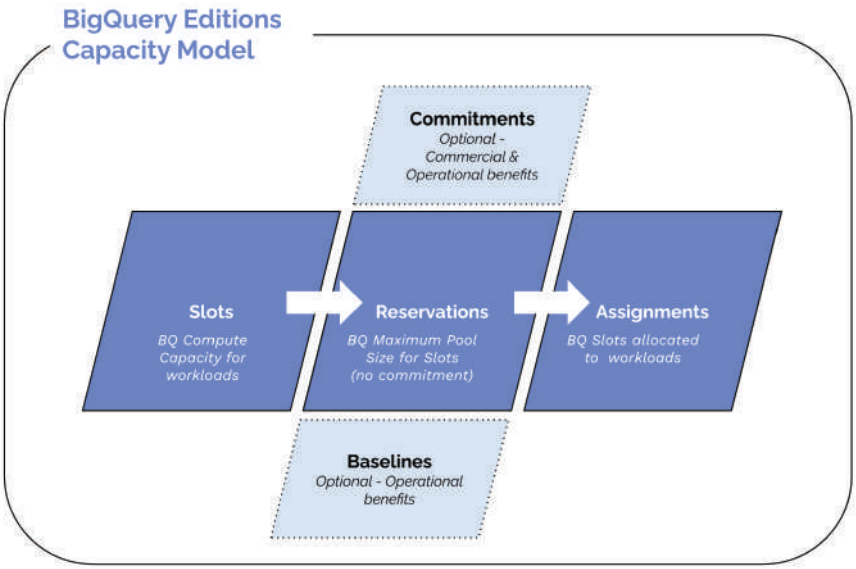

BigQuery Editions offers **autoscaling slots**, a tremendously flexible capacity planning feature. The autoscaling slots is a feature of Google's BigQuery Edition that allows users to automatically adjust the number of virtual machine (VM) slots allocated to their queries based on the size and complexity of the data being processed up to their Maximum Reservation/Assignment size.

BigQuery has allowed users to manually specify the number of slots they want to use for a query although we have very rarely seen this in use as most of the Analytics applications were not so sensitive to latency associated with BigQuery managing the slot assignments on-the-fly. BigQuery Editions autoscaling still gives us this on-the-fly capability but will require some planning with Assignments. Each level in the resource hierarchy inherits the Slot Assignment from the level above it. In other words, if a project, folder or Query is not assigned, then that query, project or folder inherits the assignment of its parent folder or organisation, if any.

BigQuery automatically adjusts the number of slots allocated to a query based on the data size and complexity, ensuring that queries are executed as quickly and efficiently as possible. The autoscaling feature works by monitoring the performance of queries in real-time and adjusting the number of slots used to ensure that queries are completed as quickly as possible. If a query is running slowly, BigQuery will automatically allocate more slots to it to speed up the processing.

With BigQuery Editions we are now obliged to define our maximum slots. This maximum is defined in the form of a Reservation which sits at the level of the GCP Organisation. The slots planning is then "soft-allocated" in the form of Assignments to the level of GCP Organisation, GCP Folder or GCP Project and subsequently the type of workload (Query, Pipeline or ML_EXTERNAL). These Assignments define the maximum number of slots based. When configuring the slot allocation in a reservation, we can additionally define a **baseline** which is the minimum number of slots dedicated to the reservation. If baseline>0 you pay for the baseline slots regardless of whether they are used. You might want to use baseline along with **committed slots.** Commitments allow the user to buy one year or three year subscription of fixed compute capacity. Committed slots are billed at a lower rate than the regular PAYG rate. The minimum commitment (and reservation) size is 100 slots, and commitments are available in 100-slot increments. There is no limit on the number of commitments that you can create. You are charged from the moment your commitment purchase is successful. Baselines are a "soft-commitment" where Commitments are "Hard". Baselines can be changed. You can increase the number of baseline slots in a reservation every few minutes. If you want to decrease your baseline slots, you are limited to once an hour if you have recently changed your baseline slot capacity and your baseline slots exceed your committed slots. Otherwise, you can decrease your baseline slots every few minutes.

An additional feature of BigQuery Editions that becomes relevant when a baseline and/or committed slots are configured is the idle slots: slots that have been assigned can be utilised by other projects in the same edition. This feature is available for Enterprise and Enterprise plus plans and only baseline and committed slots can be shared.

Our General Guidelines for Slots Capacity Planning

When performing a compute capacity planning to migrate or update your current BIgQuery on-demand or flat-rate plan to BigQuery Editions, you should consider several factors in relation with the slot consumption over a period of time.

1. The **longer** this period of time is, the **better** estimation you can provide about your resource consumption needs and forecasts. In addition, this would make the identification of potential outliers in your resource utilisation much easier to detect. Let's not forget that part of the capacity planning exercise may also be to limit active slots for given workloads to ensure a balance between cost and performance.

2. What the **maximum** and **minimum** slot consumption over this period is. Once you have spotted a maximum slot consumption value you can define the maximum number of slots for the reservation. Similarly, if the minimum slot consumption is greater than zero, you could think about defining a baseline and purchasing a commitment to save some costs. We will elaborate on this concept later.

3. What the **average** slot consumption over this period is. By combining the information of the max-/min and avg slot consumption you can have a clearer idea about the requirements of your application and decide the reservation configuration that better suits your needs.

As a rule, when defining a reservation for your application you should add your max slot consumption as the maximum slot value in your reservation. If you expect to increase the workload and still want to keep the same speed for your queries, you can opt for increasing the maximum slot limit of the reservation to up to 10-20% of your current max slot consumption (do not do this unless there is a strong reason to do so). Along with the **maximum slot value**, you can also define a **baseline** and a **commitment.** If your workload is highly peaked and the slot consumption is zero or close to zero most of the time, you do not need to set a baseline. However, if you observe that your application presents a uniform slot consumption greater than zero, you should consider purchasing slots through commitments. These slots are cheaper than regular autoscaling slots. If your uniform slot consumption is, for instance, 200 slots, you could buy 150 committed slots and set the baseline to 150. This way, you will always be charged with those 150 slots at a lower price than regular autoscaling slots and also leave some space (50 slots) to avoid incurring extra expenses in case your application workload is lower than 200 at a given time.

Overall, when setting a baseline and/or commitment be aware of the following:

BigQuery Edition Baseline Reservations This capability should be used with caution as it has a direct cost implication. We would recommend that Baselines are defined with care and only once there is a consistent and warranted pattern of consumption.

BigQuery Edition Commitments should ONLY be used once you have defined Reservation Baselines Slots and verified for a sufficient period of time that the slots are indeed being fully consumed each month. We strongly recommend a period of validation sufficient avg. out any occasional spikes in BigQuery usage. Once validated we would recommend that a Commitment be defined as a percentage of the "observed baseline slots" to ensure that the business does not over Commit, say **50-75%**.

None Assignments represent an absence of an assignment. The common use case for None assignments is to assign an organisation to the reservation and to opt out some projects or folders from that reservation by assigning them to **None**.

How to Estimate BigQuery Slot Consumption

First you will be required to have the IAM role **BigQuery Resource Viewer** (roles/bigquery.resourceViewer) to formulate your recommendations. This role is required at the Organization level for which you would like to conduct your analysis. You can also conduct the analysis at any level within the resources hierarchy for which you are conducting your analysis; Organisation, Folder or Project.

The simplest approach you can use to estimate your current slot consumption is by accessing Google's BigQuery Estimator tool in preview. This is the quickest means to get a very rapid look at your current slots consumption and get a simplified forecast of slots going forward but the greatest weakness is that it only offers a 7 day perspective of avg. and maximum slots.

From the BigQuery console you can use the BigQuery Monitoring Metrics Explorer option to view usage statistics but this view is limited to 30 days of data usage.

At Making Science we prefer to attack the challenge from a bottoms up view leveraging the data available in the INFORMATION_SCHEMA. This allows us to pick-up on any strange looking outliers and to conduct some of our own workload categorisations aligned with the customer's business and maturity of users/systems leveraging BigQuery.

At Making Science we have tools (scripts, queries and spreadsheets, etc.) to help our Data Engineers work through and analyse the usage data.

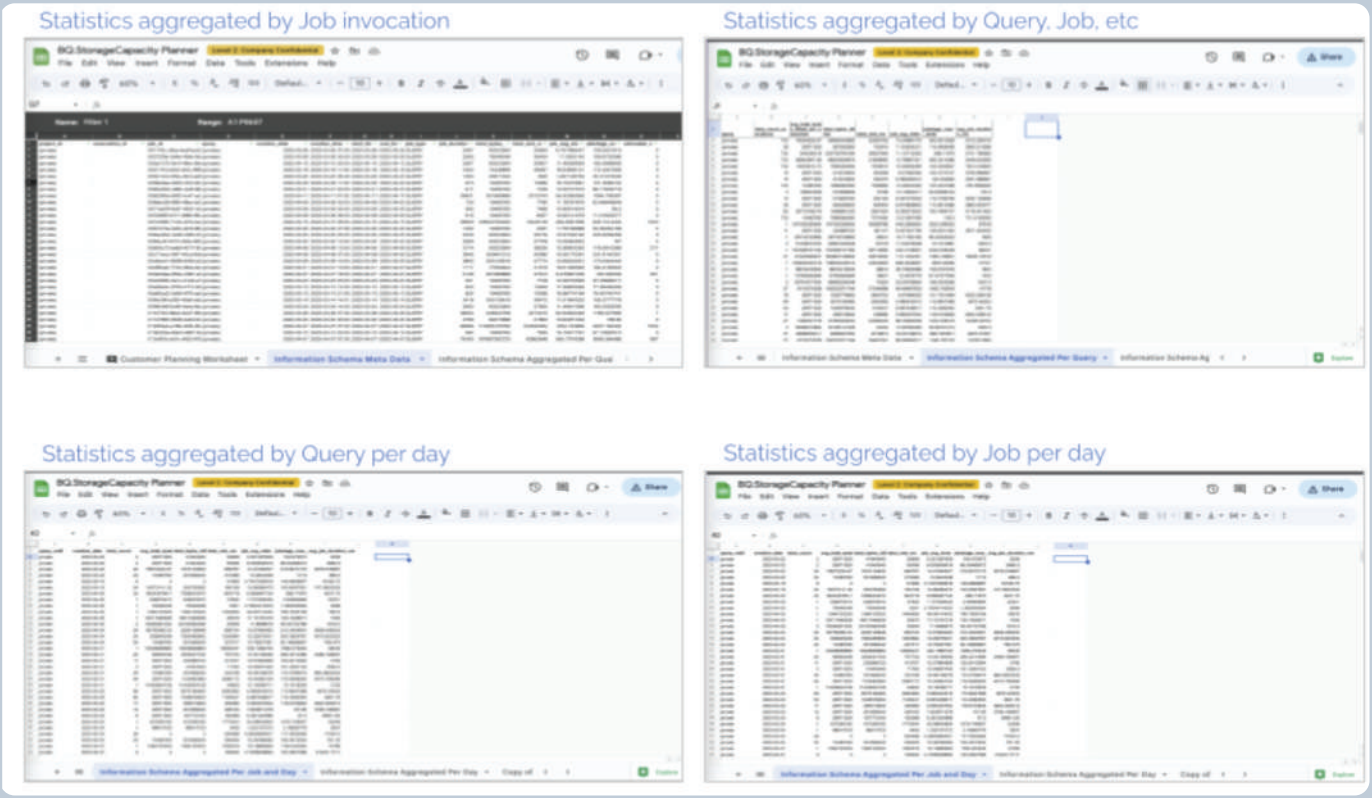

INFORMATION_SCHEMA data contains a detailed view of the slot consumption at job and also stage level that can be combined to generate insights about maximum and average slot utilisation at a given time or period. The Making Science Engineers use our compute capacity planning sheets to assist them in aggregating INFORMATION_SCHEMA data and configuring a balance of performance and costs management for BigQuery computing capacity slots. Basically the Data Engineers will hold a session with the customer to assist with the categorization phase. Attributes collected from customer interviews are; performance sensitivity of workload(s), stability of workload(s), growth forecasts for workload(s), etc. This step is paramount in optimising costs because NOT ALL JOBS/QUERIES/WORKLOADS SHOULD BE TREATED EQUALLY. This is an important premise as it forces us to evaluate the business impact "value" for having a query finish in milliseconds, seconds or even minutes.

Google recommends some very standard types of Reservations categorization groups as:
Data Pipelines, DataScience, Ad-Hoc Queries. This will help in having a more structured division of work and slot reservation when a production environment is set up.

How to Configure BigQuery Slots with BigQuery Editions

First you will be required to have the IAM role **BigQuery Resource Editor** (roles/bigquery.resourceEditor) to apply your slot recommendations to the BigQuery reservations, baselines and assignments.

When you are ready to create commitments you will need to have the IAM role
BigQuery Resource Admin (roles/bigquery.resourceAdmin)

Google recommends creating a dedicated project for BigQuery Reservations resources. This project should be named to easily identify it as the administration project, because it centralises the billing and management of any of your commitments. Give this project a descriptive name like bq-COMPANY_NAME-admin.

Remember to set the maximum number of slots according to the results of your slot capacity planning: always use the max slot consumption as the maximum slot value in your reservation. If the average slot consumption of your application is greater than zero consider also to purchase some commitments. Check Our general Guidelines for Slot Capacity Planning section for more information.

IMPORTANT NOTE: Only projects within the same Organisation resource as the administration project can be assigned to a reservation. If the administration project is not part of an Organization, no other folders or projects can use the slots.

At Making Science we follow our best practice of defining IaaS as code using Terraform. As part of any capacity planning implementation our engineers would always define the reservations, assignments and commitments using Terraform. It ensures the traceability of any changes made in GCP configurations for an Organisation.

V. Conclusions

[Google is] increasing the price of the on-demand analysis model by 25% across all regions, starting on July 5, 2023..... The vast majority of BigQuery customers around the globe utilise the On-Demand BigQuery IaaS. If you make any level of serious use of BigQuery, this will affect you.

In advance of this price increase, Google has enhanced BigQuery to give us more options with regard to controlling how we use the available features and resources. Not only will this provide vastly improved predictability of the cost, it will also provide an important opportunity for a fundamental optimisation of resource forecasting and storage usage and hence not only avoid a larger bill but potentially lower our overall spend and at the least give us the predictability with cost controls.

For larger customers, for example those spending more than $2,000.00 per month, it therefore makes sense to invest time in advance of the above date to analyse both current and planned use of BigQuery and therefrom define a reliable capacity forecast that can be used to define how we want to configure the use of slots (baselines, reservations, assignments and commitments) discussed earlier, and the appropriate utilisation of physical and logical storage.

Making Science can help you create a capacity plan, corresponding configuration recommendations and metrics to monitor their ongoing effectiveness. This exercise can typically be completed for enterprises within 2 to 8 weeks depending on the number, size and complexity of GCP projects running on BigQuery.

Our experience suggests that not all of your BigQuery projects will warrant the use of BigQuery Editions and the good news is that your smaller projects of say less than $200.00/month can remain with the BigQuery On-Demand pricing option.

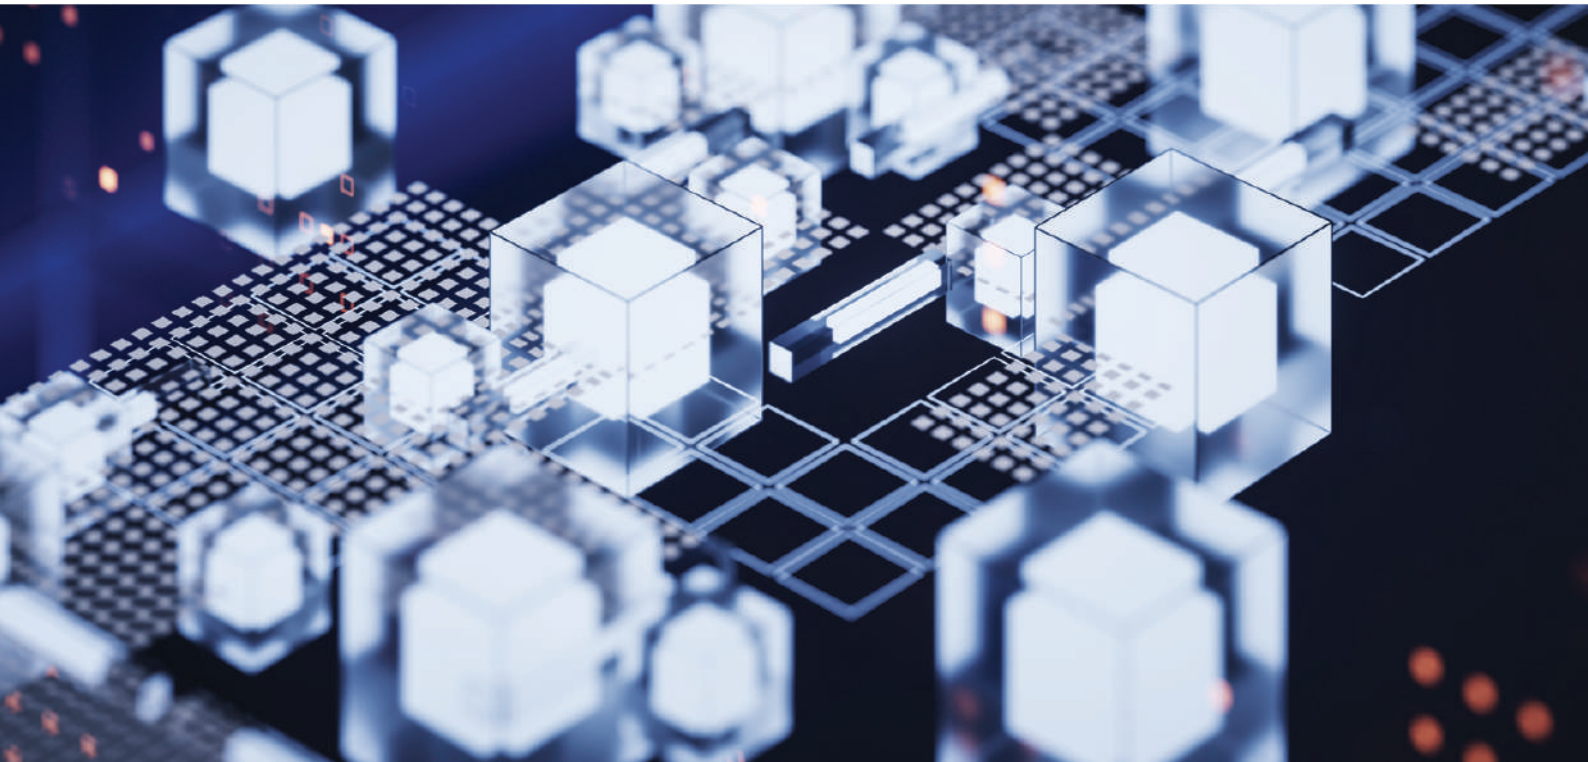

VI. References

Author	Reference
Google, G.Kazmaier VP and GM Data & Analytics GCP	New BigQuery editions: flexibility and predictability for your data cloud
Google	BigQuery Pricing Details
Google	BigQuery Editions Introduction
CMG	How to do Capacity Planning in the Cloud
Google	Introduction to BigQuery Information Schema
Google	The Big Query Estimator Tool in Preview
Google	Monitoring BigQuery Resources Utilisation
Google	Table Storage Size information
Google	Changing BigQuery Dataset to Physical Storage Billing
Google	Optimising BigQuery Slots for a Job
Google	Understanding jobs and the reservation model in BigQuery
Google	How to structure your BigQuery Resources
GitHub	A SQL framework for measuring BigQuery Slots Consumption
Google	An Insider's Guide to BigQuery Cost Optimization